

Reductions Between Expansion Problems

Prasad Raghavendra*

David Steurer†

Madhur Tulsiani‡

November 12, 2010

Abstract

The *Small-Set Expansion Hypothesis* (Raghavendra, Steurer, STOC 2010) is a natural hardness assumption concerning the problem of approximating the edge expansion of small sets in graphs. This hardness assumption is closely connected to the *Unique Games Conjecture* (Khot, STOC 2002). In particular, the Small-Set Expansion Hypothesis implies the Unique Games Conjecture (Raghavendra, Steurer, STOC 2010).

Our main result is that the Small-Set Expansion Hypothesis is in fact equivalent to a variant of the Unique Games Conjecture. More precisely, the hypothesis is equivalent to the Unique Games Conjecture restricted to instance with a fairly mild condition on the expansion of small sets. Alongside, we obtain the first strong hardness of approximation results for the BALANCED SEPARATOR and MINIMUM LINEAR ARRANGEMENT problems. Before, no such hardness was known for these problems even assuming the Unique Games Conjecture.

These results not only establish the Small-Set Expansion Hypothesis as a natural unifying hypothesis that implies the Unique Games Conjecture, all its consequences and, in addition, hardness results for other problems like BALANCED SEPARATOR and MINIMUM LINEAR ARRANGEMENT, but our results also show that the Small-Set Expansion Hypothesis problem lies at the combinatorial heart of the Unique Games Conjecture.

The key technical ingredient is a new way of exploiting the structure of the UNIQUE GAMES instances obtained from the Small-Set Expansion Hypothesis via (Raghavendra, Steurer, 2010). This additional structure allows us to modify standard reductions in a way that essentially destroys their local-gadget nature. Using this modification, we can argue about the expansion in the graphs produced by the reduction without relying on expansion properties of the underlying UNIQUE GAMES instance (which would be impossible for a local-gadget reduction).

*Georgia Institute of Technology, Atlanta, GA. Research done while visiting Princeton University

†Microsoft Research New England, Cambridge, MA. Research done at Princeton University supported by NSF Grants CCF-0832797, 0830673, and 0528414.

‡Princeton University and Institute for Advanced Study, Princeton, NJ. Work supported by NSF grant CCF-0832797 and IAS Sub-contract no. 00001583.

Contents

1	Introduction	1
1.1	Results (Informal Description)	3
2	Preliminaries	4
3	Results	5
3.1	Relation to the Unique Games Conjecture	5
3.2	Hardness Amplification for Graph Expansion	6
4	Warm-up: Hardness for Balanced Separator	7
4.1	Candidate Reduction from Unique Games	7
4.2	Structured Unique Games from Small-Set Expansion	9
4.3	Reduction from Small-Set Expansion to Balanced Separator	10
5	Additional Preliminaries	12
6	Reduction between Expansion Problems	15
6.1	Completeness	18
6.2	Soundness	20
6.2.1	Decoding a UNIQUE GAMES assignment	21
6.2.2	Decoding a very small non-expanding set in G	22
6.3	Putting things together	23
	References	24
A	Further Proofs	26
A.1	Reduction from SMALL-SET EXPANSION to UNIQUE GAMES	26
A.2	Stronger Small-Set Expansion Hypothesis	29
A.3	Hardness of MINIMUM LINEAR ARRANGEMENT and BALANCED SEPARATOR	29

1 Introduction

Finding small vertex or edge separators in a graph is a fundamental computational task. Even from a purely theoretical standpoint, the phenomenon of vertex and edge expansion – the lack of good vertex and edge separators, has had numerous implications in all branches of theoretical computer science. Yet, the computational complexity of detecting and approximating expansion, or finding good vertex and edge separators in graphs is not very well understood.

Among the two notions of expansion, this work will concern mostly with edge expansion. For simplicity, let us first consider the case of a d -regular graph $G = (V, E)$. The edge expansion of a subset of vertices $S \subseteq V$ measures the fraction of edges that leave S . Formally, the edge expansion $\Phi(S)$ of a (non-empty) subset $S \subseteq V$ is defined as,

$$\Phi_G(S) = \frac{|E(S, V \setminus S)|}{d|S|},$$

where $E(S, V \setminus S)$ denotes the set of edges with one endpoint in S and the other endpoint in $V \setminus S$. The conductance or the Cheeger’s constant associated with the graph G is the minimum of $\Phi(S)$ over all sets S with at most half the vertices, i.e.,

$$\Phi_G = \min_{|S| \leq n/2} \Phi_G(S).$$

These notions of conductance can be extended naturally to non-regular graphs, and finally to arbitrary weighted graphs (see [Section 2](#)). Henceforth, in this section, for a subset of vertices S in a graph G we will use the notation $\mu(S)$ to denote the normalized set size, i.e., $\mu(S) = |S|/n$ in a n vertex graph.

The problem of approximating the quantity Φ_G for a graph G , also referred to as the the uniform SPARSEST CUT (equivalent within a factor of 2), is among the fundamental problems in approximation algorithms. Efforts towards approximating Φ_G have led to a rich body of work with strong connections to spectral techniques and metric embeddings.

The first approximation for conductance was obtained by discrete analogues of the Cheeger inequality [[Che70](#)] shown by Alon-Milman [[AM85](#)] and Alon [[Alo86](#)]. Specifically, Cheeger’s inequality relates the conductance Φ_G to the second eigenvalue of the adjacency matrix of the graph – an efficiently computable quantity. This yields an approximation algorithm for Φ_G , one that is used heavily in practice for graph partitioning. However, the approximation for Φ_G obtained via Cheeger’s inequality is poor in terms of a approximation ratio, especially when the value of Φ_G is small. An $O(\log n)$ approximation algorithm for Φ_G was obtained by Leighton and Rao [[LR99](#)]. Later work by Linial et al. [[LLR95](#)] and Aumann and Rabani [[AR98](#)] established a strong connection between the SPARSEST CUT problem and the theory of metric spaces, in turn spurring a large and rich body of literature. More recently, in a breakthrough result Arora et al. [[ARV04](#)] obtained an $O(\sqrt{\log n})$ approximation for the problem using semidefinite programming techniques.

Small Set Expansion. It is easy to see that Φ_G is a fairly coarse measure of edge expansion, in that it is the worst case edge expansion over sets S of all sizes. In a typical graph (say a random d -regular graph), smaller sets of vertices expand to a larger extent than sets with half the vertices. For instance, all sets S of $n/1000$ vertices in a random d -regular graph have $\Phi(S) \geq 0.99$ with very high probability, while the conductance Φ_G of the entire graph is roughly $1/2$.

A more refined measure of the edge expansion of a graph is its expansion profile. Specifically, for a graph G the expansion profile is given by the curve

$$\Phi_G(\delta) = \min_{\mu(S) \leq \delta} \Phi(S) \quad \forall \delta \in [0, 1/2].$$

The problem of approximating the expansion profile has received much less attention, and is seemingly far less tractable. The second eigenvalue λ_2 fails to approximate the expansion of small sets in graphs. On one hand, even with the largest possible spectral gap, the Cheeger's inequality cannot yield a lower bound greater than $1/2$ for the conductance $\Phi_G(\delta)$. More importantly, there exists graphs such as hypercube where Φ_G is small (say ε), yet every sufficiently small set has near perfect expansion ($\Phi(S) \geq 1 - \varepsilon$). This implies that Φ_G (and the second eigenvalue λ_2) does not yield any information about expansion of small sets.

In a recent work, Raghavendra, Steurer, and Tetali [RST10] give a polynomial-time algorithm based on semidefinite programming for this problem. Roughly speaking, the approximation guarantee of their algorithm for $\Phi_G(\delta)$ is similar to the one given by Cheeger's inequality for Φ_G , except with the approximation degrading by a $\log 1/\delta$ factor. In particular, the approximation gets worse as the size of the sets considered gets smaller.

In the regime when $\Phi_G(\delta)$ tends to zero as a function of the instance size n , an $O(\log n)$ -approximation follows from the framework of Räcke [Räc08]. Very recently, this approximation has been improved to a $O(\sqrt{\log n \cdot \log(1/\delta)})$ -approximation [BKN⁺10]. Our work focuses on the regime when $\Phi_G(\delta)$ is not a function of the instance size n . In this regime, the algorithm of [RST10] gives the best known approximation for the expansion profile $\Phi_g(\delta)$.

In summary, the current state-of-the-art algorithms for approximating the expansion profile of a graph are still very far from satisfactory. Specifically, the following hypothesis is consistent with the known algorithms for approximating expansion profile.

Hypothesis (Small-Set Expansion Hypothesis, [RS10]). *For every constant $\eta > 0$, there exists sufficiently small $\delta > 0$ such that given a graph G it is NP-hard to distinguish the cases,*

Yes: *there exists a vertex set S with volume $\mu(S) = \delta$ and expansion $\Phi(S) \leq \eta$,*

No: *all vertex sets S with volume $\mu(S) = \delta$ have expansion $\Phi(S) \geq 1 - \eta$.*

For the sake of succinctness, we will refer to the above promise problem as SMALL-SET EXPANSION with parameters (η, δ) . Apart from being a natural optimization problem, the SMALL-SET EXPANSION problem is closely tied to the Unique Games Conjecture, as discussed in the next paragraph.

Recently, Arora, Barak, and Steurer [ABS10] showed that the problem SMALL-SET EXPANSION (η, δ) admits a subexponential algorithm, namely an algorithm that runs in time $\exp(n^\eta/\delta)$. However, such an algorithm does not refute the hypothesis that the problem SMALL-SET EXPANSION (η, δ) might be hard for every constant $\eta > 0$ and sufficiently small $\delta > 0$.

Unique Games Conjecture. The Khot's Unique Games Conjecture [Kho02] is among the central open problems in hardness of approximation. At the outset, the conjecture asserts that a certain constraint satisfaction problem called the Unique Games is hard to approximate in a strong sense.

An instance of UNIQUE GAMES consists of a graph with vertex set V , a finite set of labels $[R]$, and a permutation $\pi_{v \leftarrow w}$ of the label set for each edge (v, w) of the graph. A labeling $F : V \rightarrow [R]$ of the vertices of the graph is said to *satisfy* an edge (v, w) , if $\pi_{v \leftarrow w}(F(w)) = F(v)$. The objective is to find a labeling that satisfies the maximum number of edges.

The Unique Games Conjecture asserts that if the label set is large enough then even though the input instance has a labeling satisfying almost all the edges, it is NP-hard to find a labeling satisfying any non-negligible fraction of edges.

In recent years, Unique Games Conjecture has been shown to imply optimal inapproximability results for classic problems like MAX CUT [KKMO07], VERTEX COVER [KR08] SPARSEST CUT [KV05] and all constraint

satisfaction problems [Rag08]. Unfortunately, it is not known if the converse of any of these implications holds. In other words, there are no known polynomial-time *reductions* from these classic optimization problems to UNIQUE GAMES, leaving the possibility that while the its implications are true the conjecture itself could be false.

Recent work by two of the authors established a *reverse* reduction from the SMALL-SET EXPANSION problem to Unique Games [RS10]. More precisely, their work showed that Small-Set Expansion Hypothesis implies the Unique Games Conjecture. This result suggests that the problem of approximating expansion of small sets lies at the combinatorial heart of the Unique Games problem. In fact, this connection proved useful in the development of subexponential time algorithms for Unique Games by Arora, Barak and Steurer [ABS10]. It was also conjectured in [RS10] that Unique Games Conjecture is equivalent to the Small-Set Expansion Hypothesis.

1.1 Results (Informal Description)

In this work, we further investigate the connection between SMALL-SET EXPANSION and the Unique Games problem. The main result of this work is that the Small-Set Expansion Hypothesis is equivalent to a variant of the Unique Games Conjecture. More precisely, we show the following:

Theorem (Main Theorem, Informal). *The Small-Set Expansion Hypothesis is equivalent to assuming that the Unique Games Conjecture holds even when the input instances are required to be small set expanders, i.e., sets of roughly δn vertices for some small constant δ have expansion close to 1.*

As a corollary, we show that Small-Set Expansion Hypothesis implies hardness of approximation results for BALANCED SEPARATOR and MINIMUM LINEAR ARRANGEMENT problems. The significance of these results stems from two main reasons.

First, the Unique Games Conjecture is not known to imply hardness results for problems closely tied to graph expansion such as BALANCED SEPARATOR and MINIMUM LINEAR ARRANGEMENT. The reason being that the hard instances of these problems are required to have certain global structure namely expansion. Gadget reductions from a unique games instance preserve the global properties of the unique games instance such as lack of expansion. Therefore, showing hardness for BALANCED SEPARATOR or MINIMUM LINEAR ARRANGEMENT problems often required a stronger version of the Unique Games Conjecture, where the instance is guaranteed to have good expansion. To this end, several such variants of the conjecture for expanding graphs have been defined in literature, some of which turned out to be false [AKK⁺08]. Our main result shows that the Small-Set Expansion Hypothesis serves as a natural unified assumption that yields all the implications of Unique Games Conjecture and, in addition, also hardness results for other fundamental problems such as BALANCED SEPARATOR.

Second, several results in literature point to the close connection between SMALL-SET EXPANSION problem and the Unique Games problem. One of the central implications of the Unique Games Conjecture is that certain semidefinite programs yield optimal approximation for various classes of problems. As it turns out, hard instances for semidefinite programs (SDP integrality gaps) for MAX CUT [FS02, KV05, KS09, RS09], VERTEX COVER [GMPT07], UNIQUE GAMES [KV05, RS09] and SPARSEST CUT [KV05, KS09, RS09] all have near-perfect edge expansion for small sets. In case of UNIQUE GAMES, not only do all known integrality gap instances have near-perfect edge expansion of small sets, even the analysis relies directly on this property. All known integrality gap instances for semidefinite programming relaxations of Unique Games, can be translated in to gap instances for SMALL-SET EXPANSION problem, and are arguably more natural in the latter context. Furthermore, all the algorithmic results for SMALL-SET EXPANSION, including the latest work of Arora, Barak and Steurer [ABS10] extend to Unique Games as well. This apparent connection was formalized in

the result of Raghavendra et al. [RS10] which showed that Small-Set Expansion Hypothesis implies the Unique Games Conjecture. This work complements that of Raghavendra et al. [RS10] in exhibiting that the SMALL-SET EXPANSION problem lies at the combinatorial heart of the Unique Games problem.

We also show a “hardness amplification” result for SMALL-SET EXPANSION proving that if the Small-Set Expansion Hypothesis holds then the current best algorithm for SMALL-SET EXPANSION due to [RS10] is optimal within some fixed constant factor. One can view the reduction as a “scale change” operation for expansion problems, which starting from the qualitative hardness of a problem about expansion of sets with a sufficiently small measure δ , gives the optimal quantitative hardness results for problems about expansion of sets with any desired measure (larger than δ). This is analogous to (and based on) the results of [KKMO07] who gave a similar alphabet reduction for UNIQUE GAMES. An interesting feature of the reductions in the paper is that they produce instances whose expansion of small sets closely mimics a certain graph on the Gaussian space.

2 Preliminaries

Random walks on graphs. Consider the natural random walk on V defined by G . We write $j \sim G(i)$ to denote a random neighbor of vertex i in G (one step of the random walk started in i). The stationary measure for the random walk is given by the volume as defined earlier with $\mu(i) = G(\{i\}, V)$. If G is regular, then μ is the uniform distribution on V . In general, μ is proportional to the degrees of the vertices in G . We write $i \sim \mu$ to denote a vertex sampled according to the stationary measure. If G is clear from the context, we often write $i \sim V$ instead of $i \sim \mu$.

Spectral gap of graphs. We identify G with the stochastic matrix of the random walk on G . We equip the vector space $\{f: V \rightarrow \mathbb{R}\}$ with the inner product

$$\langle f, g \rangle \stackrel{\text{def}}{=} \mathbb{E}_{x \sim \mu} f(x)g(x).$$

We define $\|f\| = \langle f, f \rangle^{1/2}$. As usual, we refer to this (Hilbert) space as $L_2(V)$. Notice that G is self-adjoint with respect to this inner product, i.e., $\langle f, Gg \rangle = \langle Gf, g \rangle$ for all $f, g \in L_2(V)$. Let $\lambda_1 \geq \dots \geq \lambda_n$ be the eigenvalues of G . The non-zero constants are eigenfunctions of G with eigenvalue $\lambda_1 = 1$.

For a vertex set $S \subseteq V$, let $\mathbb{1}_S$ be the $\{0, 1\}$ -indicator function of S . We denote by $G(S, T) = \langle \mathbb{1}_S, G\mathbb{1}_T \rangle$ the total weight of all the edges in G that go between S and T .

Fact 2.1. *Suppose the second largest eigenvalue of G is λ . Then, for every function $f \in L_2(V)$,*

$$\langle f, Gf \rangle \leq (\mathbb{E} f)^2 + \lambda \cdot (\|f\|^2 - (\mathbb{E} f)^2).$$

In particular, $\Phi_G(\delta) \geq 1 - \delta - \lambda$ for every $\delta > 0$.

Gaussian Graphs. For a constant $\rho \in (-1, 1)$, let $\mathcal{G}(\rho)$ denote the infinite graph over \mathbb{R} where the weight of an edge (x, y) is the probability that two standard Gaussian random variables X, Y with correlation ρ equal x and y respectively. The expansion profile of Gaussian graphs is given by $\Phi_{\mathcal{G}(\rho)}(\mu) = 1 - \Gamma_\rho(\mu)/\mu$ where the quantity $\Gamma_\rho(\mu)$ defined as

$$\Gamma_\rho(\mu) := \mathbb{P}_{(x,y) \in \mathcal{G}_\rho} \{x \geq t, y \geq t\},$$

where \mathcal{G}_ρ is the 2-dimensional Gaussian distribution with covariance matrix

$$\begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}$$

and $t \geq 0$ is such that $\mathbb{P}_{(x,y) \sim \mathcal{G}_\rho} \{x \geq t\} = \mu$. A theorem of Borell [Bor85] shows that for any set S of measure μ , $(\mathcal{G}(\rho))(S, S) \leq \Gamma_\rho(\mu)$. This expansion profile will be frequently used in the paper to state the results succinctly.

Noise graphs. For a finite probability space (Ω, ν) and $\rho \in [0, 1]$, we define $T = T_{\rho, \Omega}$ to be the following linear operator on $L_2(\Omega)$,

$$Tf(x) = \rho x + (1 - \rho) \mathbb{E}_{x' \sim \Omega} f(x').$$

The eigenvalues of T are 1 (with multiplicity 1) and ρ (with multiplicity $|\Omega| - 1$). The operator T corresponds to the following natural (reversible) random walk on Ω : with probability ρ stay at the current position, with probability $(1 - \rho)$ move to a random position sampled according to the measure ν .

Product graphs. If G and G' are two graphs with vertex sets V and V' , we let $H = G \otimes G'$ be the *tensor product* of G and G' . The vertex set of H is $V \times V'$. For $i \in V$ and $i' \in V'$, the distribution $H(i, i')$ is the product of the distributions $G(i)$ and $G'(i')$. For $R \in \mathbb{N}$, we let $G^{\otimes R}$ denote the *R -fold tensor product* of G . Sometimes the power R of the graph is clear from the context. In this case, we might drop the superscript for the tensor graph.

3 Results

Towards stating the results succinctly, we introduce the notion of a decision problem being **SSE-hard**. It is the natural notion wherein a decision problem is **SSE-hard** if the **SMALL-SET EXPANSION** (η, δ) reduces to it by a polynomial time reduction for some constant η and all $\delta > 0$ (See Definition 5.6).

3.1 Relation to the Unique Games Conjecture

We show that the Small-Set Expansion Hypothesis is equivalent to a certain variant of the Unique Games Conjecture with expansion. Specifically, consider the following version of the conjecture with near-perfect expansion of sufficiently small sets. The hypothesis is as follows: ¹

Hypothesis 3.1 (Unique Games with Small-Set Expansion). *For every $\varepsilon, \eta > 0$ and $M \in \mathbb{N}$, there exists $\delta = \delta(\varepsilon, M) > 0$ and $q = q(\varepsilon, \eta, M) \in \mathbb{N}$ such that it is **NP-hard** to distinguish for a given **UNIQUE GAMES** instance \mathcal{U} with alphabet size q whether*

YES: *The **UNIQUE GAMES** instance \mathcal{U} is almost satisfiable, $\text{opt}(\mathcal{U}) > 1 - \varepsilon$.*

NO: *The **UNIQUE GAMES** instance \mathcal{U} satisfies $\text{opt}(\mathcal{U}) < \eta$ and its constraint graph G satisfies $\Phi(S) > 1 - \varepsilon$ for every vertex set with $\delta \leq \mu(S) \leq M\delta$.*

The main result of the paper is the following reduction from **SMALL-SET EXPANSION** to **UNIQUE GAMES** on instances with small-set expansion.

¹The hypothesis in [RS10] is not quite the same. However, the reduction and its analysis in [RS10] also work for this hypothesis.

Theorem 3.2. *For every $q \in \mathbb{N}$ and every $\varepsilon, \gamma > 0$, it is SSE-hard to distinguish between the following cases for a given UNIQUE GAMES instance \mathcal{U} with alphabet size q :*

YES: *The UNIQUE GAMES instance \mathcal{U} is almost satisfiable, $\text{opt}(\mathcal{U}) > 1 - 2\varepsilon - o(\varepsilon)$*

NO: *The optimum of the UNIQUE GAMES instance \mathcal{U} is negligible, and the expansion profile of the instance resembles the Gaussian graph $\mathcal{G}(1 - \varepsilon)$. More precisely, the UNIQUE GAMES instance \mathcal{U} satisfies $\text{opt}(\mathcal{U}) < O(q^{-\varepsilon/(2-\varepsilon)}) + \gamma$ and in addition, the constraint graph G of \mathcal{U} satisfies*

$$\forall S \subseteq V(G). \quad \Phi_G(S) \geq \Phi_{\mathcal{G}(1-\varepsilon)}(\mu(S)) - \gamma/\mu(S).$$

The proof of the above theorem is presented in [Section 6.3](#). Together with Theorem 1.9 from [\[RS10\]](#), [Theorem 3.2](#) implies the following equivalence:

Corollary 3.3. *The Small-Set Expansion Hypothesis is equivalent to [Hypothesis 3.1](#) (Unique Games with Small-Set Expansion).*

Remark 3.4. If we choose $\gamma \ll \varepsilon$, then the constraint graph G in the No case satisfies $\Phi(S) \geq \Omega(\sqrt{\varepsilon})$ for every vertex set S with $\mu(S) \in (b, 1/2)$ for an arbitrarily small constant $b > 0$. In other words, the best balanced separator in G has cost $\Omega(\sqrt{\varepsilon})$. A hardness of Unique Games on graphs of this nature was previously conjectured in [\[AKK⁺08\]](#), towards obtaining a hardness for BALANCED SEPARATOR.

As already mentioned, for several problems such as MAX CUT, the hard instances for the semidefinite programs have very good expansion of small sets. For instance, hard instances for semidefinite programs (SDP integrality gaps) for MAX CUT [\[FS02, KV05, KS09, RS09\]](#), VERTEX COVER [\[GMPT07\]](#), UNIQUE GAMES [\[KV05, RS09\]](#) and SPARSEST CUT [\[KV05, KS09, RS09\]](#) all have near-perfect edge expansion for small sets. In fact, in many of the cases, the edge expansion in the graph closely mimics the expansion of sets in some corresponding Gaussian graph. Confirming this observation, our techniques imply an optimal hardness result for MAX CUT on instances that are small-set expanders. More precisely, the Small-Set Expansion Hypothesis implies that the Goemans-Williamson algorithm is optimal even on graphs that are guaranteed to have good expansion of small sets, in fact an expansion profile that resembles the Gaussian graph. For the sake of succinctness, we omit the formal statement of the result.

3.2 Hardness Amplification for Graph Expansion

Observe that the Small-Set Expansion Hypothesis is a purely qualitative assumption on the approximability of expansion. Specifically, for every constant η the hypothesis asserts that there exists some δ such that approximating expansion of sets of size δ is NP-hard. The hypothesis does not assert any quantitative dependence on the set size and approximability. Surprisingly, we show that this qualitative hardness assumption is sufficient to imply precise quantitative bounds on approximability of graph expansion.

Theorem 3.5. *For all $q \in \mathbb{N}$ and $\varepsilon, \gamma > 0$, it is SSE-hard to distinguish between the following two cases for a given graph $H = (V_H, E_H)$*

YES: *There exist q disjoint sets $S_1, \dots, S_q \subseteq V_H$ satisfying for all $l \in [q]$,*

$$\mu(S_l) = \frac{1}{q} \quad \text{and} \quad \Phi_H(S_l) \leq \varepsilon + o(\varepsilon).$$

No: For all sets $S \subseteq V_H$,

$$\Phi_H(S) \geq \Phi_{\mathcal{G}(1-\varepsilon/2)}(\mu(S)) - \gamma/\mu(S)$$

where $\Phi_{\mathcal{G}(1-\varepsilon/2)}(\mu(S))$ is the expansion of sets of volume $\mu(S)$ in the infinite Gaussian graph $\mathcal{G}(1-\varepsilon/2)$.

The above hardness result matches (upto an absolute constant factor), the recent algorithmic result (Theorem 1.2) of [RST10] approximating the graph expansion. Furthermore, both the YES and the No cases of the above theorem are even qualitatively stronger than in the Small-Set Expansion Hypothesis. In the YES case, not only does the graph have one non-expanding set, but it can be partitioned into small sets, *all* of which are non-expanding. This partition property is useful in some applications such as hardness reduction to MINIMUM LINEAR ARRANGEMENT. In the No case, the expansion of all sets can be characterized only by their size $\mu(S)$. Specifically, the expansion of every set S of vertices with $\mu(S) \gg \gamma$, is at least the expansion of a set of similar size in the Gaussian graph $\mathcal{G}(1-\varepsilon/2)$.

Here we wish to draw an analogy to the Unique Games Conjecture. The Unique Games Conjecture is qualitative in that it does not prescribe a relation between its soundness and alphabet size. However, the work of Khot et al. [KKMO07] showed that the Unique Games Conjecture implies a quantitative form of itself with a precise relation between the alphabet size and soundness. Theorem 3.5 could be thought of as an analogue of this phenomena for the Small-Set Expansion problem.

As an immediate consequence of Theorem 3.5, we obtain the following hardness of the BALANCED SEPARATOR and MINIMUM LINEAR ARRANGEMENT problems (See Appendix A.3 for details).

Corollary 3.6 (Hardness of BALANCED SEPARATOR and MIN BISECTION). *There is a constant c such that for arbitrarily small $\varepsilon > 0$, it is SSE-hard to distinguish the following two cases for a given graph $G = (V, E)$:*

YES: *There exists a cut $(S, V \setminus S)$ in G such that $\mu(S) = \frac{1}{2}$ and $\Phi_G(S) \leq \varepsilon + o(\varepsilon)$.*

No: *Every cut $(S, V \setminus S)$ in G , with $\mu(S) \in (\frac{1}{10}, \frac{1}{2})$ satisfies $\Phi_G(S) \geq c\sqrt{\varepsilon}$.*

Corollary 3.7 (Hardness of MINIMUM LINEAR ARRANGEMENT). *It is SSE-hard to approximate MINIMUM LINEAR ARRANGEMENT to any fixed constant factor.*

4 Warm-up: Hardness for Balanced Separator

In this section we present a simplified version of our reduction from SMALL-SET EXPANSION to BALANCED SEPARATOR. Though it gives sub-optimal parameters, it illustrates the key ideas used in the general reduction.

4.1 Candidate Reduction from Unique Games

A natural approach for reducing UNIQUE GAMES to BALANCED SEPARATOR is to consider variants of the reduction from UNIQUE GAMES to MAX CUT in [KKMO07] (similarly, one could consider variants of the reduction from UNIQUE GAMES to the *generalized* SPARSEST CUT problem [KV05]).

Let \mathcal{U} be a unique game with alphabet size R and vertex set V . (We assume that every vertex of the unique game participates in the same number of constraints. This assumption is without loss of generality.) The candidate reduction has a parameter $\varepsilon > 0$. The graph $H = H_\varepsilon(\mathcal{U})$ obtained from this candidate reduction has vertex set $V \times \{0, 1\}^R$ and its edge distribution is defined as follows:

1. Sample a random vertex $u \in V$.

2. Sample two random constraints $(u, v, \pi), (u, v', \pi')$ of \mathcal{U} that contain the vertex u . (Henceforth, we will write $(u, v, \pi) \sim \mathcal{U} \mid u$ to denote a random constraint of \mathcal{U} containing vertex u .)
3. Sample a random edge (y, y') of the boolean noise graph $T_{1-\varepsilon}$ with noise parameter ε .
4. Output an edge between $(v, \pi(y))$ and $(v', \pi'(y'))$. (Here, $\pi(y)$ denotes the vector obtained by permuting the coordinates of y according to the permutation π .)

Completeness. Suppose there is a good assignment $F: V \rightarrow [R]$ for the unique game \mathcal{U} . Then, if we sample a random vertex $u \in V$ and two random constraint $(u, v, \pi), (u, v', \pi') \sim \mathcal{U} \mid u$, with probability very close to 1 (much closer than ε), the labels assigned to v and v' satisfy $\pi^{-1}(F(v)) = (\pi')^{-1}(F(v'))$. Consider the vertex set $S = \{(u, x) \mid x_{F(u)} = 1\}$ in the graph H . We have $\mu(S) = 1/2$. We claim that the expansion of this set is essentially $\varepsilon/2$ (up to a lower-order term depending on the fraction of constraint of \mathcal{U} violated by F). Consider a random edge e with endpoints $(v, \pi(y))$ and $(v', \pi'(y'))$, where the vertices $v, v' \in V$ and the permutations π, π' are generated as specified above. Let $r = \pi^{-1}(F(v))$ and $r' = (\pi')^{-1}(F(v'))$. The edge e crosses the cut S if and only if $y_r \neq y'_{r'}$. As argued before, with probability very close to 1, we have $r = r'$. Conditioned on this event, the probability that $y_r \neq y'_{r'}$ is equal to $\varepsilon/2$. This shows that S has expansion $\varepsilon/2$.

Soundness. Suppose no assignment for the unique game \mathcal{U} satisfies a significant fraction of constraints. Let S be a vertex set in the graph H . The goal is to lower bound the expansion of S (which is the same as upper bounding the fraction of edges with both endpoints in S). Let $f: V^R \times \{0, 1\}^R \rightarrow \{0, 1\}$ be the indicator function of S . Following the analysis of [KKMO07], we consider functions $g_u: \{0, 1\}^R \rightarrow [0, 1]$,

$$g_u(x) = \mathbb{P}_{\substack{(u,v,\pi) \sim \mathcal{U} \mid u, \\ y \sim T_{\sqrt{1-\varepsilon}}(x)}} \{f(v, \pi(y))\}.$$

(The graph H turns out to be the square of a graph H_0 in which we would just create edges of the form $((u, x), (v, \pi(y)))$. The function $g_u(x)$ evaluates to the probability that the set S contains a random neighbor of (u, x) in this graph H_0 .) By construction, the fraction $H(S, S)$ of edges of H with both endpoints in S is exactly

$$H(S, S) = \mathbb{E}_{u \in V} \langle g_u, T_{1-\varepsilon} g_u \rangle.$$

Since \mathcal{U} does not have a good assignment, standard arguments (invariance principle and influence decoding, see [KKMO07]) imply the following upper bound on $H(S, S)$,

$$H(S, S) \leq \mathbb{E}_{u \in V} \Gamma_{1-\varepsilon}(\mu_u) + o(1).$$

(The notation $o(1)$ hides a term depending on the maximum fraction of constraints of \mathcal{U} that can be satisfied. For us, this term is not significant.) Here, μ_u is the expected value of g_u and $\Gamma_{1-\varepsilon}(\cdot)$ is the noise stability profile of the Gaussian noise graph with parameter ε . We would like to show that every set S that contains a μ fraction of the vertices of H satisfies $H(S, S) \leq \Gamma_{1-\varepsilon}(\mu) + o(1)$. However, the function $\Gamma_{1-\varepsilon}$ is, of course, not concave. Hence, this upper bound holds only if μ_u is close to μ for most vertices $u \in V$.

In fact, it is very easy to construct examples that show that the candidate reduction is not sound. For example, consider a unique game \mathcal{U} that consists of two disjoint parts of the same size (i.e., without any constraint between the two parts). The reduction preserves this global structure, in the sense that the graph H also consists of two disjoint parts of the same size (with no edge between the parts). Hence, this graph contains a vertex set with volume $1/2$ and expansion 0 irrespective of the optimal value of the unique game \mathcal{U} .

In fact, any cut in the underlying graph of \mathcal{U} can be translated to a cut in H and the resulting function f may have the values μ_u as (very close to) 0 or 1.

This example shows that the above candidate reduction can only work if one makes assumptions about structure of the constraint graph of the underlying unique game \mathcal{U} . However, such an assumption raises the question if UNIQUE GAMES could be hard to approximate even if the constraint graph is expanding. This issue turns out to be delicate as demonstrated by the algorithm for UNIQUE GAMES with expanding constraint graphs [AKK⁺08]. This algorithm achieves a good approximation for UNIQUE GAMES if the expansion of the constraint graph exceeds a certain threshold.

4.2 Structured Unique Games from Small-Set Expansion

In this work, we present a very different approach for fixing the above candidate reduction. Instead of assuming expansion properties of the constraint graph, we assume that the underlying unique game is obtained by the reduction from SMALL-SET EXPANSION to UNIQUE GAMES in [RS10]². This specific form of the underlying unique game will allow us to modify the reduction such that the global structure of the constraint graph is no longer preserved in the graph obtained from the reduction. (In particular, our modified reduction will break with the paradigm of composing unique games with local gadgets.)

In the following, we describe the reduction from SMALL-SET EXPANSION to UNIQUE GAMES. Let G be a regular graph with vertex set V . For technical reasons, we assume that G contains a copy of the complete graph of weight $\eta > 0$. (Since we will be able to work with very small η , this assumption is without loss of generality.) Given a parameter $R \in \mathbb{N}$ and the graph G , the reduction outputs a unique game $\mathcal{U} = \mathcal{U}_R(G)$ with vertex set V^R and alphabet $[R]$. The constraints of the unique game \mathcal{U} correspond to the following probabilistic verifier for an assignment $F: V^R \rightarrow [R]$:

1. Sample a random vertex $A \in V^R$.
2. Sample two random neighbors $B, C \sim G^{\otimes R}(A)$ of the vertex A in the tensor-product graph $G^{\otimes R}$.
3. Sample two random permutations π_B, π_C of $[R]$.
4. Verify that $\pi_B^{-1}(F(\pi_B(B))) = (\pi_C)^{-1}(F(\pi_C(C)))$.

Raghavendra and Steurer [RS10] show that this reduction is complete and sound in the following sense:

Completeness If the graph G contains a vertex set with volume $1/R$ and expansion close to 0, then the unique game $\mathcal{U} = \mathcal{U}_R(G)$ has a partial assignment that labels an $\alpha \geq 1/e$ fraction of the vertices and satisfies almost an α fraction of the constraints.

Soundness If the graph G contains no set with volume $1/R$ and expansion bounded away from 1, then no assignment for the unique game $\mathcal{U} = \mathcal{U}_R(G)$ satisfies a significant fraction of the constraints.

Hence, if one assumes the Small-Set Expansion Hypothesis, then the kind of unique games obtained from the reduction are hard to approximate.

We remark that the completeness of the reduction seems weaker than usual, because we are only guaranteed a partial assignment for the unique game. However, it is easy to check that the KKMO reduction presented in the previous section also works if there is only a partial assignment in the completeness case. The only difference is that one now gets a set S with $\mu(S) = \alpha/2$ and expansion roughly $\varepsilon/2$.

²We remark that unique games of this form do not necessarily have expanding constraint graphs. In fact, it is still possible that the constraint graph consists of two disconnected components.

4.3 Reduction from Small-Set Expansion to Balanced Separator

We now show how the combination of the above two reductions can be modified to give a reduction from SMALL-SET EXPANSION to BALANCED SEPARATOR. Let $\mathcal{U} = \mathcal{U}_R(G)$ be the unique game given by the reduction of Raghavendra and Steurer. If we consider the graph H_ε given by the reduction in Section 4.1, each vertex of H_ε is now of the form (A, x) , where $A \in V^R$ and $x \in \{0, 1\}^R$.

The intuition is that in this case, we can think of x as picking a *subset* of the vertices in A , and that just the knowledge of this subset (instead of the whole of A) is sufficient for the provers to provide a good answer to the corresponding unique game. In particular, let $A' = \{A_i \mid x_i = 1\}$ is the subset picked by x . Then the argument for the completeness case in [RS10] actually shows that one can still find a good labeling for an α fraction of the vertices A , where the label of A only depends on A' ³.

Formally, if we replace A with the tuple $A'(x)$ defined by taking $A'_i = A_i$ if $x_i = 1$ and $A'_i = \perp$ otherwise. This gives a graph H' with the vertex set being a subset of $(V \cup \{\perp\})^R \times \{0, 1\}^R$. The the argument in completeness case for showing that H has a balanced cut of expansion roughly $\varepsilon/2$ can in fact be extended to show that H' also has a balanced cut of expansion roughly $\varepsilon/2$.

The soundness analysis in the previous reduction did not always work because H had the same structure as $G^{\otimes R}$, since we essentially replaced every vertex of $G^{\otimes R}$ by a gadget $\{0, 1\}^R$ to obtain H . However, the structure of H' is very different from that of $G^{\otimes R}$.

For example, consider the vertices $A = (u_1, u_2, \dots, u_R)$ and $B = (v_1, u_2, \dots, u_R)$ in V^R which only differ in the first coordinate (A, B are not necessarily adjacent). Let $x \in \{0, 1\}^R$ be such that $x_1 = 0$. Then, while (A, x) and (B, x) are different vertices in H , $(A'(x), x)$ and $(B'(x), x)$ are in fact the same vertex in H' ! On the other hand, if $x_1 = 1$, then $(A'(x), x)$ and $(B'(x), x)$ would be two different vertices in H' . Hence, the gadget structure of H is no longer preserved in H' - it is very different from a “locally modified” copy of $G^{\otimes R}$.

For the purposes of analysis, it will be more convenient to think of A' being obtained by replacing A_i where $x_i = 0$, by a random vertex of G instead of the symbol \perp . Instead of identifying different vertices in H with the same vertex in H' , this now has the effect of re-distributing the weight of an incident on (A, x) , uniformly over all the vertices that (A', x) can map to. Let M_x denote a Markov operator which maps A to a random A' as above (a more general version and analysis of such operators can be found in Section 5).

We now state the combined reduction. The weight of an edge in the final graph H' is the probability that it is produced by the following process:

1. Sample a random vertex $A \in V^R$.
2. Sample two random neighbors $B, C \sim G^{\otimes R}(A)$ of the vertex A in the tensor-product graph $G^{\otimes R}$.
3. Sample $x_B, x_C \sim \{0, 1\}^R$.
4. Sample $B' \sim M_{x_B}(B)$ and $C' \sim M_{x_C}(C)$.
5. Sample two random permutations π_B, π_C of $[R]$.
6. Output an edge between the vertices $\pi_B(B', x_B)$ and $\pi_C(C', x_C)$ ($\pi(A, x)$ denotes the tuple $(\pi(A), \pi(x))$).

As before, let $f : V^R \times \{0, 1\}^R$ denote the indicator function of a set in H' , with (say) $\mathbb{E} f = \mu = 1/2$. We define the functions

$$\tilde{f}(A, x) \stackrel{\text{def}}{=} \mathbb{E}_\pi f(\pi.A, \pi.x) \quad \text{and} \quad g_A(x) \stackrel{\text{def}}{=} \mathbb{E}_{B \sim G^{\otimes R}(A)} \mathbb{E}_{B' \sim M_x(B)} \tilde{f}(B', x).$$

³Given a non-expanding small set S , if $A' \cap S$ contains a single element A'_j , then we assign the label j to A . If $A' \cap S$ is not a singleton, we do not label A .

By construction, each vertex (A, x) of H' has exactly the same neighborhood structure as $(\pi.A', \pi.x)$ for all $\pi \in S_R$ and $A' \in M_x(A)$. Hence, the fraction of edges crossing the cut can also be written in terms of \bar{f} as $\langle f, H'f \rangle = \langle \bar{f}, H'\bar{f} \rangle$.

We will show that \bar{f} gives a cut (actually, a distribution over cuts) with the same expansion in the graph H , such that the functions g_A satisfy $\mathbb{P}_A \{ \mathbb{E} g_A \in (1/10, 9/10) \} \geq 1/10$. Recall that showing this was exactly the problem in making the reduction in [Section 4.1](#) work.

Since $\mathbb{E}_A \mathbb{E}_x g_A = \mu$, we have $\mathbb{E}_A (\mathbb{E}_x g_A)^2 \geq \mu^2$. The following claim also gives an upper bound.

Claim 4.1. $\mathbb{E}_A (\mathbb{E}_x g_A)^2 \leq \mu^2/2 + \mu/2$

Proof. We have

$$\begin{aligned} \mathbb{E}_{A \sim V^R} \left(\mathbb{E}_x g_A \right)^2 &= \mathbb{E}_{A \sim V^R} \left(\mathbb{E}_{B \sim G^{\otimes R}(A)} \mathbb{E}_x \mathbb{E}_{B' \sim M_x(B)} \bar{f} \right)^2 \leq \mathbb{E}_{A \sim V^R} \mathbb{E}_{B \sim G^{\otimes R}} \left(\mathbb{E}_x \mathbb{E}_{B' \sim M_x(B)} \bar{f} \right)^2 \\ &= \mathbb{E}_{B \sim V^R} \left(\mathbb{E}_x \mathbb{E}_{B' \sim M_x(B)} \bar{f} \right)^2 \\ &= \mathbb{E}_{B \sim V^R} \left[\left(\mathbb{E}_{x_1 B'_1 \sim M_{x_1}(B)} \bar{f} \right) \left(\mathbb{E}_{x_2 B'_2 \sim M_{x_2}(B)} \bar{f} \right) \right] \\ &= \mathbb{E}_{x_1 B'_1 \sim M_{x_1}(B)} \bar{f}(B'_1, x_1) \mathbb{E}_{(B'_2, x_2) \sim M(B'_1, x_1)} f(B'_2, x_2). \end{aligned}$$

For the last equality above, we define M to be a Markov operator which samples (B'_2, x_2) from the correct distribution given (B'_1, x_1) . Since x_1, x_2 are independent, x_2 can just be sampled uniformly. The fact that B'_1 and B'_2 come from the same (random) B can be captured by sampling each coordinate of B'_2 as

$$(B'_2)_i = \begin{cases} (B'_1)_i & \text{if } (x_1)_i = (x_2)_i = 0 \\ \text{random vertex in } V & \text{otherwise} \end{cases}.$$

Abusing notation, we also use M to denote the operator on the space of the functions which averages the value of the function over random (B'_2, x_2) generated as above. Then, if λ is the second eigenvalue of M , we have

$$\mathbb{E}_A \left(\mathbb{E}_x g_A \right)^2 \leq \langle \bar{f}, M\bar{f} \rangle \leq 1 \cdot (\mathbb{E} \bar{f})^2 + \lambda \cdot (\|\bar{f}\|^2 - (\mathbb{E} \bar{f})^2) \leq (1 - \lambda) \cdot \mu^2 + \lambda \cdot \mu.$$

Finally, it can be checked that the second eigenvalue of M is $1/2$ which proves the claim. \square

This gives that $\mathbb{E}_x g_A$ cannot be always very far from μ . Formally,

$$\mathbb{P}_A \{ |\mathbb{E} g_A - \mu| \geq \gamma \} \leq \frac{\mathbb{E}_A (\mathbb{E} g_A - \mu)^2}{\gamma^2} \leq \frac{\mu(1 - \mu)}{2\gamma^2}.$$

Hence, for $\gamma = 2/5$, the probability is at most $25/32 < 9/10$. This can now be combined with the bound from [Section 4.1](#) that gives

$$H'(S, S) \leq \mathbb{E}_A \Gamma_{1-\varepsilon}(\mathbb{E} g_A) + o(1).$$

Since $\mathbb{E} g_A \geq 1/10$ with probability at least $1/10$ over A , these “nice” A ’s contribute a volume of at least $1/100$. Also, for a nice A , we have $\Gamma_{1-\varepsilon}(\mathbb{E} g_A) \leq (\mathbb{E} g_A)(1 - \Omega(\sqrt{\varepsilon}))$. Hence,

$$H'(S, S) \leq (\mu - 1/100) + 1/100 \cdot (1 - \Omega(\sqrt{\varepsilon})) + o(1)$$

which shows that S has expansion $\Omega(\sqrt{\varepsilon})$.

5 Additional Preliminaries

Unique Games. An instance of UNIQUE GAMES represented as $\mathcal{U} = (\mathcal{V}, \mathcal{E}, \Pi, [R])$ consists of a graph over vertex set \mathcal{V} with the edges \mathcal{E} between them. Also part of the instance is a set of labels $[R] = \{1, \dots, R\}$, and a set of permutations $\Pi = \{\pi_{v \leftarrow w} : [R] \rightarrow [R]\}$, one permutation for each edge $e = (w, v) \in \mathcal{E}$. An assignment $F : \mathcal{V} \rightarrow [R]$ of labels to vertices is said to satisfy an edge $e = (w, v)$, if $\pi_{v \leftarrow w}(F) = F(v)$. The objective is to find an assignment F of labels that satisfies the maximum number of edges.

As is customary in hardness of approximation, one defines a gap-version of the UNIQUE GAMES problem as follows:

Problem 5.1 (UNIQUE GAMES $(R, 1 - \varepsilon, \eta)$). Given a UNIQUE GAMES instance $\mathcal{U} = (\mathcal{V}, \mathcal{E}, \Pi = \{\pi_{v \leftarrow w} : [R] \rightarrow [R] \mid e = (w, v) \in \mathcal{E}\}, [R])$ with number of labels R , distinguish between the following two cases:

- $(1 - \varepsilon)$ -satisfiable instances: There exists an assignment F of labels that satisfies a $1 - \varepsilon$ fraction of edges.
- Instances that are not η -satisfiable: No assignment satisfies more than a η -fraction of the edges \mathcal{E} .

The Unique Games Conjecture asserts that the above decision problem is NP-hard when the number of labels is large enough. Formally,

Conjecture 5.2 (Unique Games Conjecture [Kho02]). *For all constants $\varepsilon, \eta > 0$, there exists large enough constant R such that UNIQUE GAMES $(R, 1 - \varepsilon, \eta)$ is NP-hard.*

Graph expansion. In this work, all graphs are undirected and possibly weighted. Let G be a graph with vertex set V . We write $ij \sim G$ to denote a random edge sampled from G (with random orientation). For two vertex sets $S, T \subseteq V$, let $G(S, T)$ be the fraction of edges going from S to T , i.e.,

$$G(S, T) \stackrel{\text{def}}{=} \mathbb{P}_{ij \sim G} \{i \in S, j \in T\}.$$

The *expansion*⁴ $\Phi_G(S)$ of a set $S \subseteq V$ is the fraction of edges leaving S normalized by the fraction of edges incident to S , i.e.,

$$\Phi_G(S) \stackrel{\text{def}}{=} \frac{G(S, V \setminus S)}{G(S, V)} = \frac{\mathbb{P}_{ij \sim G} \{i \in S, j \notin S\}}{\mathbb{P}_{ij \sim G} \{i \in S\}} = \mathbb{P}_{ij \sim G} \{j \notin S \mid i \in S\}.$$

The *volume* of a set S is the fraction of edges incident on it and is denoted by $\mu(S) \stackrel{\text{def}}{=} G(S, V)$. The fraction of edges leaving the set is denoted by $\partial(S) \stackrel{\text{def}}{=} G(S, V \setminus S)$.

Small-Set Expansion Hypothesis.

Problem 5.3 (SMALL-SET EXPANSION (η, δ)). Given a regular graph $G = (V, E)$, distinguish between the following two cases:

YES: There exists a non-expanding set $S \subseteq V$ with $\mu(S) = \delta$ and $\Phi_G(S) \leq \eta$.

No: All sets $S \subseteq V$ with $\mu(S) = \delta$ are highly expanding having $\Phi_G(S) \geq 1 - \eta$.

⁴The technically more precise term is *conductance*

Hypothesis 5.4 (Hardness of approximating SMALL-SET EXPANSION). *For all $\eta > 0$, there exists $\delta > 0$ such that the promise problem SMALL-SET EXPANSION (η, δ) is NP-hard.*

Remark 5.5. It is easy to see that for the problem SMALL-SET EXPANSION (η, δ) to be hard, one must have $\delta \leq \eta$. This follows from the fact that if we randomly sample a set S containing a δ fraction of the vertices (and hence, having volume δ for a regular graph), the expected fraction of edges crossing the set is $\delta(1 - \delta)$ and hence $\mathbb{E} \Phi_G(S) = 1 - \delta$. However, for it to be possible that for all sets with $\mu(S) = \delta$ have $\Phi_G(S) \geq 1 - \eta$, we must have $\delta \leq \eta$.

Definition 5.6. Let \mathcal{P} be a decision problem of distinguishing between two disjoint families (cases) of instances denoted by {Yes, No}. For a given instance I of \mathcal{P} , let $\text{Case}(I)$ denote the family to which I belongs. We say that \mathcal{P} is SSE-hard if for some $\eta > 0$ and all $\delta \in (0, \eta)$, there is a polynomial time reduction, which starting from an instance $G = (V, E)$ of SMALL-SET EXPANSION (η, δ) , produces an instance I of \mathcal{P} such that

- $\exists S \subseteq V$ with $\mu(S) = \delta$ and $\Phi_G(S) \leq \eta \implies \text{Case}(I) = \text{Yes}.$
- $\forall S \subseteq V$ with $\mu(S) = \delta$, $\Phi_G(S) \geq 1 - \eta \implies \text{Case}(I) = \text{No}.$

For the proofs, it shall be more convenient to use the following version of the SMALL-SET EXPANSION problem, in which we high expansion is guaranteed not only for sets of measure δ , but also within an arbitrary multiplicative factor of δ .

Problem 5.7 (SMALL-SET EXPANSION (η, δ, M)). Given a regular graph $G = (V, E)$, distinguish between the following two cases:

Yes: There exists a non-expanding set $S \subseteq V$ with $\mu(S) = \delta$ and $\Phi_G(S) \leq \eta$.

No: All sets $S \subseteq V$ with $\mu(S) \in \left(\frac{\delta}{M}, M\delta\right)$ have $\Phi_G(S) \geq 1 - \eta$.

The following proposition shows that for the purposes of showing that \mathcal{P} is SSE-hard, it is sufficient to give a reduction from SMALL-SET EXPANSION (η, δ, M) for any chosen values of η, M and for all δ . We defer the proof to [Section A.2](#).

Proposition 5.8. *For all $\eta > 0, M \geq 1$ and all $\delta < 1/M$, there is polynomial time reduction from SMALL-SET EXPANSION $(\frac{\eta}{M}, \delta)$ to SMALL-SET EXPANSION (η, δ, M) .*

Invariance principle. The following theorem on the noise stability of functions over a product probability space is an easy corollary of Theorem 4.4 in Mossel et al. [\[MOO05\]](#). Recall that $\Gamma_\rho(\mu) := \mathbb{P}_{(x,y) \sim \mathcal{G}_\rho} \{x \geq t, y \geq t\}$, where \mathcal{G}_ρ is the 2-dimensional Gaussian distribution with covariance matrix $\begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}$ and $t \geq 0$ is such that $\mathbb{P}_{(x,y) \sim \mathcal{G}_\rho} \{x \geq t\} = \mu$.

Theorem 5.9. *Let $\nu > 0, \rho \in (0, 1)$ and let Ω be a finite probability space. Then, there exists $\tau, \delta > 0$ such that the following holds: Every function $f: \Omega^R \rightarrow [0, 1]$ satisfies either*

$$\langle f, T_\rho f \rangle \leq \Gamma_\rho(\mathbb{E} f) + \nu.$$

or $\max_{i \in [R]} \inf_i (T_{1-\delta} f) > \tau$. (Here, T_ρ and $T_{1-\delta}$ are the natural noise operators on $L_2(\Omega^R)$ with correlation parameters ρ and $1 - \delta$ as defined above.)

Below we define generalizations of the operators M_x and M used in [Section 4](#). We show that these operators can be viewed as somewhat extended versions of the noise operators which randomize each coordinate of a product space with some probability. The operators we define can be viewed as noise operators with additional “leakage” property, in the sense that part of the output encodes the information about which coordinates were randomized. The second eigenvalue of these operators can be easily estimated by relating it to the eigenvalue of the corresponding noise operator.

Random walks with leaked randomness. Suppose we have a collection of graphs $\{G_z\}_{z \in \mathcal{Z}}$ with the same vertex set V (and with the same stationary distribution). We consider two (reversible) random walks defined by this collection and compare their spectral properties. The first random walk is defined on V . If the current state is x^1 , we choose the next state x^2 by sampling a random index $z \sim \mathcal{Z}$ and then taking two random steps from x^1 in G_z , i.e., we sample $x \sim G_z(x^1)$ and $x^2 \sim G_z(x)$. The second random walk is defined on $V \times \mathcal{Z}$. If the current state is (x^1, z^1) , we choose the next state (x^2, z^2) by sampling a random neighbor x of x^1 in G_{z^1} , then we choose a random index $z^2 \sim \mathcal{Z}$ and a random neighbor $x^2 \sim G_{z^2}(x)$ according to G_{z^2} . The following lemma shows that these two random walks have the same non-zero eigenvalues. (Recall that we identify graphs with their stochastic operators.)

Lemma 5.10. *Let (\mathcal{Z}, ν) be a finite probability space and let $\{G_z\}_{z \in \mathcal{Z}}$ be a family of graphs with the same vertex set V and stationary measure μ . Then the following two graphs have the same non-zero eigenvalues:*

- the graph $\mathbb{E}_{z \sim \mathcal{Z}} G_z^2$ on V ,
- the graph H on $V \times \mathcal{Z}$ defined by

$$Hf(x^1, z^1) = \mathbb{E}_{x \sim G_{z^1}(x^1)} \mathbb{E}_{z^2 \sim \mathcal{Z}} \mathbb{E}_{x^2 \sim G_{z^2}(x)} f(x^2, z^2).$$

Proof. Let M be the following linear operator on $L_2(V \times \mathcal{Z})$,

$$Mf(x, z) = \mathbb{E}_{z' \sim \mathcal{Z}} \mathbb{E}_{x' \sim G_{z'}(x)} f(x', z').$$

Notice that its adjoint operator M^* (with respect to the inner product in $L_2(V \times \mathcal{Z})$) is given by

$$M^*f(x, z) = \mathbb{E}_{x' \sim G_z(x)} \mathbb{E}_{z' \sim \mathcal{Z}} f(x', z').$$

(The operator above is the adjoint of M , because each of the random walks G_z are reversible and have the same stationary measure.). The graph H corresponds to the operator M^*M , which has the same non-zero eigenvalues as MM^* . The operator MM^* is given by

$$MM^*f(x^1, z^1) = \mathbb{E}_{z \sim \mathcal{Z}} \mathbb{E}_{x \sim G_z(x^1)} \mathbb{E}_{x^2 \sim G_z(x)} \mathbb{E}_{z^2 \sim \mathcal{Z}} f(x^2, z^2).$$

The subspace $\{f \mid \forall x \in V. \mathbb{E}_z f(x, z) = 0\} \subseteq L_2(V \times \mathcal{Z})$ is part of the kernel of MM^* . Hence, all eigenfunctions with non-zero eigenvalue are in the orthogonal complement of this space. The orthogonal complement consists of all functions f such that $f(x, z)$ does not depend on z . Let f be such a function and set $f(x) = f(x, z)$. Then,

$$MM^*f(x^1) = \mathbb{E}_{z \sim \mathcal{Z}} \mathbb{E}_{x \sim G_z(x^1)} \mathbb{E}_{x^2 \sim G_z(x)} f(x^2) = \mathbb{E}_{z \sim \mathcal{Z}} G_z^2 f(x^1).$$

Thus, MM^* acts on this subspace in the same way as $\mathbb{E}_z G_z^2$, which means that the two operators have the same eigenfunctions (and eigenvalues) in this space. \square

Noise graph with leaked randomness. Let $\{\perp, \top\}_\beta^R$ be the β -biased R -dimensional boolean hypercube. If we sample a random point z from this space, then $z_i = \top$ with probability β , independently for each coordinate $i \in [R]$.

Let (Ω, ν) be a finite probability space. For $z \in \{\perp, \top\}_\beta^R$ and $x \in \Omega^R$, let $M_z(x)$ be the distribution over Ω^R obtained by “rerandomizing” every coordinate of x where z has value \perp . In order to sample $x' \sim M_z(x)$, we sample $x'_i \sim \Omega$, independently for every coordinate $i \in [R]$ with $z_i = \perp$. If $z_i = \top$, then we copy the value of x in this coordinate so that $x'_i = x_i$. Observe that $\mathbb{E}_{z \sim \{\perp, \top\}_\beta^R} M_z = T_{\beta, \Omega}^{\otimes R}$ is the usual noise graph on Ω^R with correlation parameter β , as defined previously in this section.

Consider the following stochastic linear operator M on $L_2(\Omega^R, \{\perp, \top\}_\beta^R)$,

$$Mf(x, z) = \mathbb{E}_{z' \sim \{\perp, \top\}_\beta^R} \mathbb{E}_{x' \sim M_z(x)} f(x', z'). \quad (5.1)$$

The following lemma shows that the second largest singular value of M is the same as the second largest eigenvalue of the corresponding noise graph.

Lemma 5.11. *Let $f \in L_2(\Omega^R, \{\perp, \top\}_\beta^R)$ and let M be as in (5.1). Then,*

$$\|Mf\|^2 \leq (\mathbb{E} f)^2 + \beta \cdot (\|f\|^2 - (\mathbb{E} f)^2).$$

Proof. We have $\|Mf\|^2 = \langle f, M^* M f \rangle$ where M^* is the adjoint of M . This operator $M^* M$ is the same as the (second) operator in Lemma 5.10 for $G_z = M_z$. Hence, $M^* M$ has the same non-zero eigenvalues as $\mathbb{E}_z M_z^2$. From the definition of M_z , it is clear that $M_z^2 = M_z$. Further, $T = \mathbb{E}_z M_z$ is the noise operator on Ω^R with correlation parameter β . We conclude that $M^* M$ has second largest eigenvalue β . The lemma follows from Fact 2.1. \square

6 Reduction between Expansion Problems

Let G be a graph with vertex set V and stationary measure μ . Our reduction maps G to a graph H with vertex set $V^R \times \Omega^R$ for $\Omega = [q] \times \{\perp, \top\}_\beta$. Here, $R, q \in \mathbb{N}$ and $\beta > 0$ are parameters of the reduction. We impose the natural product measure on Ω ,

$$\mathbb{P}((\alpha, z)) = \begin{cases} \frac{\beta}{1-\beta} & \text{if } z = \top \\ \frac{q}{q} & \text{if } z = \perp \end{cases} \quad \forall \alpha \in [q]$$

As before, we describe H in terms of a probabilistic process defined by G , which generates the edge distribution of H . (See Figure 1 for a more condensed description.) The process uses the following three auxiliary graphs (already introduced in §2 and §5):

- First, the noise graph $T_V := T_{1-\varepsilon_V, V}^{\otimes R}$, which resamples independently every coordinate of a given R -tuple $A \in V^R$ with probability ε_V . (Here, $\varepsilon_V > 0$ is again a parameter of the reduction. We think of ε_V as rather small compared to other parameters.) This noise effectively adds a copy of the complete graph with weight ε_V to G , which we assumed in §4.2.
- Next, the noise graph $T_\Omega := T_{\rho, \Omega}^{\otimes R}$, which resamples independently every coordinate of a given R -tuple $(x, z) \in \Omega^R$ with probability $1 - \rho$. (For $x \in [q]^R$ and $z \in \{\perp, \top\}_\beta^R$, we write $(x, z) \in \Omega^R$ to denote the tuple obtained by merging corresponding coordinates of x and z to an element of Ω . In other words,

we identify $[q]^R \times \{\perp, \top\}_\beta^R$ and Ω^R .) The correlation parameter ρ of T_Ω is the most important parameter of the reduction, because the graph T_Ω plays the role of a dictatorship test gadget in our reduction. We think of ρ as being close to 1.

- Finally, we consider the graph M_z on $\tilde{\Omega}^R$ for $\tilde{\Omega} = V \times [q]$ and $z \in \{\perp, \top\}_\beta^R$. For $(A, x) \in \tilde{\Omega}^R$, the graph M_z resamples every coordinate in which z has value \perp .

Our reduction proceeds in three phases:

In the first phase, we sample a random vertex $A \in V^R$ and take two independent random steps from A according to the graph $T_V G^{\otimes R}$, i.e., we sample \tilde{B} and \tilde{C} from the distribution $T_V G^{\otimes R}(A)$. We end the first phase by sampling two random permutations π_B and π_C . The permutations are required to satisfy the property that if we divide the domain $[R]$ into contiguous blocks of size R/k , then each such block is permuted in place. We define the set Π_k of such permutations as

$$\Pi_k := \{\pi \in S_R \mid \forall j \in \{0, \dots, k-1\}. \pi(\{jR/k + 1, \dots, (j+1)R/k\}) = \{jR/k + 1, \dots, (j+1)R/k\}\}.$$

This phase exactly corresponds to the reduction from SMALL-SET EXPANSION to UNIQUE GAMES in [RS10].

In the second phase, we sample a random R -tuple (x_A, z_A) in Ω^R and take two independent random steps from (x_A, z_A) according to the graph T_Ω , i.e., we sample (x_B, z_B) and (x_C, z_C) from $T_\Omega(x_A, z_A)$. This phase corresponds to typical dictatorship test reduction (as in [KKMO07]).

In the third phase, we apply the graphs M_{z_B} and M_{z_C} to the R -tuples (\tilde{B}, x_B) and (\tilde{C}, x_C) respectively, to obtain (B', x'_B) and (C', x'_C) . The final step of this phase is to output an edge between $\pi_B(B', x'_B, z_B)$ and $\pi_C(C', x'_C, z_C)$. (For a permutation π of $[R]$ and an R -tuple X , we denote by $\pi(X)$ the permutation of X according to π , so that $(\pi(X))_{\pi(i)} = X_i$.)

We remark that the random permutations π_B and π_C in the first phase and the resampling according to M_z in the third phase introduce symmetries in the graph H that effectively identify vertices. In particular, any two vertices in $V^R \times \Omega^R$ of the form (A, x, z) and $\pi(A, x, z)$ have the same neighbors in H (i.e., the distributions $H(A, x, z)$ and $H(\pi(A, x, z))$ are identical). This kind of symmetry has been used in integrality gap constructions (see [KV05]) and hardness reductions (see [RS10]).

The kind of symmetry introduced by the M_z graph in the third phase seems to be new. In the third phase, we effectively identify vertices (A, x, z) and (A', x', z) if they differ only in the coordinates in which z has value \perp . Formally, the vertex (A, x, z) has the same distribution of neighbors as the vertex (A', x', z) if (A', x') is sampled from $M_z(A, x)$.

Remark 6.1 (Reduction to UNIQUE GAMES with expansion). We note that the above reduction can also be viewed as creating a UNIQUE GAMES instance with alphabet size q . For a vertex $(A, x, z) \in V_H$ and $l \in [q]$, let $(A, x, z) + l$ denote the vertex (A, x', z) , where $x'_i \equiv x_i + l \pmod q$ for all $i \in [R]$. We define an equivalence relation on V_H by taking $(A, x, z) \equiv (A, x, z) + l$ for all A, x, z and $l \in [q]$. Let $H/[q]$ be a graph with one vertex for each equivalence class of the above relation. Also, for each edge in E_H , we add an edge in $H/[q]$ between the equivalence classes containing the corresponding vertices of V_H . We claim that $H/[q]$ can then be viewed as a UNIQUE GAMES instance⁵ as described in Theorem 3.2.

We now describe the constraints for the edges in $H/[q]$. We identify each equivalence class with an arbitrarily chosen representative element in it. For $(A, x, z) \in V_H$, let $[(A, x, z)]$ denote the representative of the equivalence class containing it. Consider an edge in E_H between $(\pi_B(B', x'_B, z_B))$ and $(\pi_C(C', x'_C, z_C))$.

⁵In the terminology used in the literature, one can say that the graph H is a *label-extended* graph of a UNIQUE GAMES instance with alphabet size $[q]$.

The Reduction

Input: A weighted graph G with vertex set V .

Parameters: $R, q, k \in \mathbb{N}$, and $\varepsilon_V, \beta, \rho > 0$.

Output: A graph $H = (V_H, E_H)$ with vertex set $V_H = V^R \times [q]^R \times \{\top, \perp\}_\beta^R$.

Let Π_k denote the set of permutations of $[R]$ which permute each block of size R/k in-place i.e.

$$\Pi_k := \{\pi \in S_R \mid \forall j \in \{0, \dots, k-1\}. \pi(\{jR/k + 1, \dots, (j+1)R/k\}) = \{jR/k + 1, \dots, (j+1)R/k\}\}.$$

The weight of an edge in E_H is proportional to the probability that the following probabilistic process outputs this edge:

– Reducing from SMALL-SET EXPANSION to UNIQUE GAMES.

1. Sample an R -tuple of vertices $A \sim V^R$.
2. Sample two random neighbors $B, C \sim G^{\otimes R}(A)$ of A .
3. Sample $\tilde{B} \sim T_V(B)$ and $\tilde{C} \sim T_V(C)$.
4. Sample two permutations $\pi_B, \pi_C \in \Pi_k$

– Combination with long code gadgets.

6. Sample $(x_A, z_A) \in \Omega^R$, where $\Omega = [q] \times \{\perp, \top\}_\beta$.
7. Sample $(x_B, z_B), (x_C, z_C) \sim T_\Omega(x_A, z_A)$.

– Redistributing the edge weights

7. Sample $(B', x'_B) \sim M_{z_B}(\tilde{B}, x_B)$ and $(C', x'_C) \sim M_{z_C}(\tilde{C}, x_C)$
8. Output an edge between $(\pi_B(B', x'_B, z_B))$ and $(\pi_C(C', x'_C, z_C))$.

Figure 1: Reduction between expansion problems

Let $\pi_B(B', x'_B, z_B) = [\pi_B(B', x'_B, z_B)] + l_B$ and $\pi_C(C', x'_C, z_C) = [\pi_C(C', x'_C, z_C)] + l_C$. Then the constraint corresponding to this edge requires that an assignment mapping vertices in $H/[q]$ to $[q]$ must satisfy

$$F([\pi_B(B', x'_B, z_B)]) + l_B \equiv F([\pi_C(C', x'_C, z_C)]) + l_C \pmod{q}.$$

We note that the expansion properties of H are inherited by $H/[q]$, since any set of measure μ in $H/[q]$ is also a set of measure μ in H . In the YES case, each of the sets S_1, \dots, S_q mentioned in [Theorem 3.5](#) will provide an assignment for the above UNIQUE GAMES instance, satisfying $1 - \varepsilon - o(\varepsilon)$ fraction of the constraints. In the No case, we will argue that each assignment corresponds to a set of measure $1/q$ in H , and the unsatisfiability of the instance will follow from the expansion of the corresponding sets in H .

6.1 Completeness

Lemma 6.2. *Let $H = (V_H, E_H)$ be constructed from $G = (V, E)$ as in the reduction in Figure 1. If there is a set $S \subseteq V$ satisfying $\mu(S) \in \left[\frac{k}{10\beta R}, \frac{k}{\beta R}\right]$ and $\Phi_G(S) \leq \eta$, then there exists a partition S_1, \dots, S_q of V_H satisfying:*

1. *For all $(A, x, z) \in V_H$ and $l, l' \in [q]$, $(A, x, z) \in S_l \implies (A, x, z) + l' \in S_{l+l'}$.*
2. *For each all $l \in [q]$, $\Phi_H(S_l) \leq 2(1 - \rho^2 + \eta + 2\varepsilon_V) + (1 - \rho^2 + \eta + 2\varepsilon_V)^2 + \frac{(1-\rho^2)\beta}{\rho^2} + 2^{-\Omega(k)}$.*

Note that the first property, together with the fact that S_1, \dots, S_q form a partition also implies that for all $l \in [q]$, $\mu(S_l) = \frac{1}{q}$.

Proof. We first describe a procedure for assigning vertices in V_H to S_1, \dots, S_q . This procedure assigns all but $2^{-\Omega(k)}$ fraction of the vertices, which we shall distribute arbitrarily later. Let (A, x, z) be a vertex in V_H , where $A \in V^R$, $x \in [q]^R$ and $z \in \{\top, \perp\}^R$.

For all $j \in [k]$, we define the sets $W_j := \{i \in \{(j-1)R/k + 1, \dots, jR/k\} \mid z_i \neq \perp\}$. Let $A(W_j)$ denote the multiset $A(W_j) = \{A_i \mid i \in W_j\}$. We take,

$$j^* = \inf \left\{ j \mid |A(W_j) \cap S| = 1 \right\}$$

If $|A(W_j) \cap S| \neq 1$ for any $j \in [k]$, then we do not assign the vertex (A, x, z) to any of the set S_1, \dots, S_q . Else, let A_{i^*} be the unique element in $A(W_{j^*}) \cap S$. We assign

$$(A, x, z) \in S_{x_{i^*}}.$$

Note that the assignment to sets is determined only by the coordinates $i \in [R]$ for which $z_i \neq \perp$. The first property is easily seen to be satisfied for all the vertices that are assigned, as the sets $\{W_j\}_{j \in [k]}$ are identical for (A, x, z) and $(A, x, z) + l$, for any $l \in [q]$. The following claim proves that most vertices are indeed assigned to one of the sets S_1, \dots, S_q .

Claim 6.3. $\mathbb{P}_{(A, x, z) \sim V_H} \left\{ |A(W_j) \cap S| \neq 1 \ \forall j \in [k] \right\} \leq 2^{-\Omega(k)}.$

Proof. Note that over the choice of a random $(A, x, z) \in V_H$, the intersection sizes $|W_1 \cap S|, \dots, |W_k \cap S|$ are independent random variables distributed as $\text{Binomial}(\mu\beta, R/k)$. The probability that all of them are not equal to 1, can then be bounded as

$$\begin{aligned} \mathbb{P}_{(A, x, z) \sim V_H} \left\{ |A(W_j) \cap S| \neq 1 \ \forall j \in [k] \right\} &= \left(1 - \frac{R}{k} \cdot \mu\beta \cdot (1 - \mu\beta)^{R/k-1} \right)^k \\ &\leq \left(1 - \frac{R}{k} \cdot \frac{k}{10R} \cdot (1 - k/R)^{R/k} \right)^k \leq \left(1 - \frac{1}{30} \right)^k. \end{aligned}$$

The last inequality assumes that $R/k \geq 4$ so that $(1 - k/R)^{R/k} > 1/3$. □

We now bound the expansion of these sets. A random edge is between two tuples of the form $(\pi_B(B', x'_B, z_B))$ and $(\pi_C(C', x'_C, z_C))$, where π_B and π_C are two random permutations in Π_k and B', C' are generated from $G^{\otimes R}$ as in Figure 1. For a fixed $l \in [q]$, the expansion of S_l is equal to the following probability taken over the choice of a random edge

$$\mathbb{P} \left\{ (\pi_C(C', x'_C, z_C)) \notin S_l \mid (\pi_B(B', x'_B, z_B)) \in S_l \right\} = \mathbb{P} \left\{ (C', x'_C, z_C) \notin S_l \mid (B', x'_B, z_B) \in S_l \right\}.$$

Here we used the fact that the membership in a set S_l is invariant under permutations in Π_k . The following claim analyzes the above probability.

Claim 6.4. $\mathbb{P}\left\{\left(C', x'_C, z_C\right) \notin S_l \mid \left(B', x'_B, z_B\right) \in S_l\right\} \leq 2(1 - \rho^2 + 2\varepsilon_V + \eta) + (1 - \rho^2 + 2\varepsilon_V + \eta)^2 + \frac{(1 - \rho^2)\beta}{\rho^2}.$

Proof. Let $\{W_j^{(B)}\}_{j \in [k]}$ denote the multisets $W_j^{(B)} = \{i \in \{(j-1)k/R+1, \dots, jk/R\} \mid (z_B)_i \neq \perp\}$ and let $\{W_j^{(C)}\}_{j \in [k]}$ be defined similarly. Define $j_B^* = \inf\{j \mid |B'(W_j^{(B)}) \cap S| = 1\}$ when the set on the right is non-empty and $k+1$ otherwise. Let j_C^* be the analogous quantity for C' . In the cases when $j_B^*, j_C^* \leq k$, let $W_{j_B^*}^{(B)} \cap S = \{B'_{i_B^*}\}$ and $W_{j_C^*}^{(C)} \cap S = \{C'_{i_C^*}\}$. We can bound the required probability by the probability that either $j_B^* \neq j_C^*$ or $i_B^* \neq i_C^*$ or $(x'_C)_{i_C^*} \neq l$.

$$\begin{aligned} \mathbb{P}\left\{\left(C', x'_C, z_C\right) \notin S_l \mid \left(B', x'_B, z_B\right) \in S_l\right\} &\leq \mathbb{P}\left\{j_C^* \neq j_B^* \mid \left(B', x'_B, z_B\right) \in S_l\right\} \\ &\quad + \mathbb{P}\left\{\left(j_C^* = j_B^*\right) \wedge \left(i_C^* \neq i_B^*\right) \mid \left(B', x'_B, z_B\right) \in S_l\right\} \\ &\quad + \mathbb{P}\left\{\left(j_C^* = j_B^*\right) \wedge \left(i_C^* = i_B^*\right) \wedge \left((x'_C)_{i_C^*} \neq l\right) \mid \left(B', x'_B, z_B\right) \in S_l\right\} \\ &\leq \mathbb{P}\left\{j_C^* \neq j_B^* \mid j_B^* \leq k\right\} + \mathbb{P}\left\{i_C^* \neq i_B^* \mid \left(j_B^* = j_C^*\right) \wedge \left(j_B^* \leq k\right)\right\} \\ &\quad + \mathbb{P}\left\{\left(x'_C\right)_{i_C^*} \neq l \mid \left((x'_B)_{i_B^*} = l\right) \wedge \left(i_B^* = i_C^*\right)\right\} \end{aligned}$$

In the second inequality above, we drop conditionings that are irrelevant and use $\mathbb{P}\{A \wedge B\} \leq \mathbb{P}\{A \mid B\}$. We now analyze each of the above terms separately.

The first term can be further split as

$$\mathbb{P}\left\{j_C^* \neq j_B^* \mid j_B^* \leq k\right\} = \mathbb{P}\left\{j_C^* > j_B^* \mid j_B^* \leq k\right\} + \mathbb{P}\left\{j_C^* < j_B^* \mid j_B^* \leq k\right\}.$$

To have $j_C^* > j_B^*$, it must be the case that $|W_{j_B^*}^{(C)} \cap S| \neq 1$, while we also have $|W_{j_B^*}^{(B)} \cap S| = 1$ by definition of j_B^* . If $i_B^* \notin W_{j_B^*}^{(C)}$, this must be because $(z_C)_{i_B^*} = \perp$ or $C'_{i_B^*} \notin S$. The former happens with probability $1 - \rho^2$ as we already have that $(z_B)_{i_B^*} = \top$. The latter even happens with probability at most $\eta + 2\varepsilon_V$ as it could be due to the edge $(B'_{i_B^*}, C'_{i_B^*})$ going out of S or one of the vertices being perturbed by T_V . Combining, we get a bound of $(1 - \rho^2 + \eta + 2\varepsilon_V)$ for the case when i' such that $C'_{i'} \in S$ and the events above must happen for $W_{j_B^*}^{(B)}$ and i' , giving again a bound of $(1 - \rho^2 + \eta + 2\varepsilon_V)$. The term $\mathbb{P}\left\{j_C^* < j_B^* \mid j_B^* \leq k\right\}$ can be bound identically. We then get

$$\mathbb{P}\left\{j_C^* \neq j_B^* \mid j_B^* \leq k\right\} = 2(1 - \rho^2 + \eta + 2\varepsilon_V).$$

We now consider the term $\mathbb{P}\left\{i_C^* \neq i_B^* \mid \left(j_B^* = j_C^*\right) \wedge \left(j_B^* \leq k\right)\right\}$. For this to happen, the above events must occur for *both* the pairs $\left(i_B^*, W_{j_C^*}^{(C)}\right)$ and $\left(i_C^*, W_{j_B^*}^{(B)}\right)$. This gives

$$\mathbb{P}\left\{i_C^* \neq i_B^* \mid \left(j_B^* = j_C^*\right) \wedge \left(j_B^* \leq k\right)\right\} \leq (1 - \rho^2 + \eta + 2\varepsilon_V)^2$$

Finally, given $j_B^* = j_C^*$ and $i_B^* = i_C^*$, the probability that $(x'_B)_{i_B^*} \neq (x'_C)_{i_C^*}$ is at most $\frac{1 - \rho^2}{\rho^2}\beta$. This is true because for any i , $((x'_B)_i, (z_B)_i)$ and $((x'_C)_i, (z_C)_i)$ are same with probability ρ^2 and uniform in Ω with

probability $1 - \rho^2$. Also, for any i , $i = i_B^* = i_C^*$ in particular means that $(z_B)_i = (z_C)_i = \top$. Conditioned on this, we can bound the probability $(x'_B)_{i_B^*} \neq (x'_C)_{i_C^*}$ as

$$\mathbb{P}\{(x'_B)_i \neq (x'_C)_i \mid (z_B)_i = (z_C)_i = \top\} = \frac{(1 - \rho^2) \cdot (1 - 1/q) \cdot \beta^2}{(1 - \rho^2) \cdot \beta^2 + \rho^2 \cdot \beta} \leq \frac{(1 - \rho^2)\beta}{\rho^2}$$

Combining the bounds for the three terms proves the claim. \square

It remains to partition the vertices not assigned to any of the sets S_1, \dots, S_q . We simply assign any such vertex (A, x, z) to the set S_{x_1} . It is easy to see that S_1, \dots, S_q still satisfy the first property. Since the measure of the extra vertices added to each set is $\frac{2^{-\Omega(k)}}{q}$, the expansion of each set increases by at most $2^{-\Omega(k)}$. \square

6.2 Soundness

Let G be a graph with vertex set V and stationary measure μ . Let H be the graph obtained from the reduction in Figure 1. The vertex set of H is $V^R \times \Omega^R$. Recall that $\Omega = [q] \times \{\perp, \top\}_\beta$. Let $f: V^R \times \Omega^R \rightarrow [0, 1]$. We think of f as a cut in H (or convex combination thereof).

We define two symmetrizations of f as follows

$$\bar{f}(A, x, z) = \mathbb{E}_{\pi \sim \Pi_k} f(\pi(A, x, z)) \quad \text{and} \quad \bar{f}'(A, x, z) = \mathbb{E}_{(A', x') \sim M_z(A, x)} \bar{f}(\pi(A', x', z)).$$

By the symmetries of the graph,

$$\langle f, Hf \rangle = \langle \bar{f}, H\bar{f} \rangle = \langle \bar{f}', H\bar{f}' \rangle.$$

We write $\bar{f}'_A(x, z) = \bar{f}'(A, x, z)$ and consider the average (with noise) of \bar{f}'_B over the neighbors B of a vertex A in G^R ,

$$g_A = \mathbb{E}_{B \sim G^R(A)} \mathbb{E}_{\tilde{B} \sim T_V(B)} \bar{f}'_{\tilde{B}}.$$

We will apply the techniques of [KKMO07] to analyze the functions g_A . We first express the fraction of edges that stay within the cut defined by f in terms of the functions g_A .

Lemma 6.5.

$$\langle f, Hf \rangle = \mathbb{E}_{A \sim V^R} \|T_\Omega g_A\|^2.$$

Proof. Using the construction of H and the symmetry of \bar{f}' , we get

$$\begin{aligned} \langle f, Hf \rangle &= \mathbb{E}_{A \sim V^R} \mathbb{E}_{(x, z) \sim \Omega^R} \left(\mathbb{E}_{B \sim G^R(A)} \mathbb{E}_{\tilde{B} \sim T_V(B)} \mathbb{E}_{(x_B, z_B) \sim T_\Omega(x, z)} \bar{f}'_B(x_B, z_B) \right)^2 \\ &= \mathbb{E}_{A \sim V^R} \mathbb{E}_{(x, z) \sim \Omega^R} \left(\mathbb{E}_{B \sim G^R(A)} \mathbb{E}_{\tilde{B} \sim T_V(B)} T_\Omega \bar{f}'_B(x, z) \right)^2 \\ &= \mathbb{E}_{A \sim V^R} \mathbb{E}_{(x, z) \sim \Omega^R} \left(T_\Omega g_A(x, z) \right)^2 \\ &= \mathbb{E}_{A \sim V^R} \|T_\Omega g_A\|^2. \end{aligned} \quad \square$$

We now show that for most tuples A , the functions g_A have the same expectation as $\mathbb{E} f$. To this end, we show that $\mathbb{E}_A (\mathbb{E} g_A)^2 \approx (\mathbb{E} f)^2$.

Lemma 6.6.

$$\mathbb{E}_{A \sim V^R} \left(\mathbb{E}_{\Omega^R} g_A \right)^2 \leq (\mathbb{E} f)^2 + \beta \|f\|^2.$$

Proof. Let $\tilde{\Omega} = V \times [q]$. We have

$$\begin{aligned} \mathbb{E}_{A \sim V^R} \left(\mathbb{E}_{\Omega^R} g_A \right)^2 &= \mathbb{E}_{A \sim V^R} \left(\mathbb{E}_{B \sim G^R(A)} \mathbb{E}_{\tilde{B} \sim T_V(B)} \mathbb{E}_{(x,z) \sim \Omega^R} \tilde{f}'(B, x, z) \right)^2 \\ &\leq \mathbb{E}_{A \sim V^R} \mathbb{E}_{B \sim G^R(A)} \mathbb{E}_{\tilde{B} \sim T_V(B)} \mathbb{E}_x \left(\mathbb{E}_z \tilde{f}'(B, x, z) \right)^2 \quad (\text{Cauchy-Schwarz}) \\ &= \mathbb{E}_{(A,x) \sim \tilde{\Omega}^R} \left(\mathbb{E}_z \tilde{f}'(A, x, z) \right)^2. \end{aligned} \quad (6.1)$$

Let M be the following stochastic operator on $L_2(\tilde{\Omega}^R \times \{\perp, \top\}_\beta^R)$,

$$Mh(A, x, z_0) = \mathbb{E}_{z \sim \{\perp, \top\}_\beta^R} \mathbb{E}_{(A', x') \sim M_z(A, x)} h(A', x', z).$$

Recall that $\tilde{f}'(A, x, z) = \mathbb{E}_{(A', x') \sim M_z(A, x)} \tilde{f}(A', x', z)$. With this notation, the right-hand side of (6.1) simplifies to $\|M\tilde{f}\|^2$. Therefore,

$$\mathbb{E}_{A \sim V^R} \left(\mathbb{E}_{\Omega^R} g_A \right)^2 \leq \|M\tilde{f}\|^2 \leq (\mathbb{E} \tilde{f})^2 + \beta \|\tilde{f}\|^2 \leq (\mathbb{E} f)^2 + \beta \|f\|^2.$$

The second inequality uses that M^*M has second largest eigenvalue β (see Lemma 5.11). The last inequality uses that \tilde{f} is obtained by applying a stochastic operator on f . \square

The following lemma is an immediate consequence of the previous lemma (Lemma 6.6) and Chebyshev's inequality.

Lemma 6.7. *For every $\gamma > 0$,*

$$\mathbb{P}_{A \sim V^R} \left\{ \mathbb{E} g_A \geq \mathbb{E} f + \gamma \sqrt{\mathbb{E} f} \right\} \leq \beta/\gamma^2 \cdot \frac{\|f\|^2}{\mathbb{E} f} \leq \beta/\gamma^2.$$

Proof. Lemma 6.6 shows that $\mathbb{E}_A(g_A - \mathbb{E} f)^2 \leq \beta \|f\|^2$. Hence, $\mathbb{P}_A \left\{ |g_A - \mathbb{E} f| > \gamma \sqrt{\mathbb{E} f} \right\} \leq \beta \|f\|^2 / (\gamma^2 \mathbb{E} f)$. \square

6.2.1 Decoding a UNIQUE GAMES assignment

The goal is decode from f an assignment $F: V^R \rightarrow [R]$ that maximizes the probability

$$\mathbb{E}_{A \sim V^R, B \sim G^R(A)} \mathbb{E}_{\tilde{A} \sim T_V(A)} \mathbb{E}_{\tilde{B} \sim T_V(B)} \mathbb{P}_{\pi_A, \pi_B \sim \Pi_k} \left\{ \pi_A^{-1}(F(\pi(A))) = \pi_B^{-1}(F(\pi(B))) \right\}. \quad (6.2)$$

(The expression above is roughly the success probability of the assignment F for the UNIQUE GAMES instance obtained by applying the reduction from [RS10] on G .)

As usual, we decode according to influential coordinates of f (after symmetrization). More precisely, we generate a assignment F by the following probabilistic process: For every $A \in V^R$, with probability $1/2$, choose a random coordinate in $\{i \in [R] \mid \text{Inf}_i(T_{1-\delta} g_A) > \tau\}$ and with probability $1/2$, choose a random coordinate in $\{i \in [R] \mid \text{Inf}_i(T_{1-\delta} \tilde{f}'_A) > \tau\}$. If the sets of influential coordinates are empty, we choose a uniformly random coordinate in $[R]$.

The following lemma follows immediately from the techniques in [KKMO07]. The reason is that (6.2) is the success probability of the assignment F for a UNIQUE GAMES instance defined on V^R . For $A \in V^R$, the function g_A is the average over bounded functions $f'_B: \Omega^R \rightarrow [0, 1]$, where B is a random neighbor of A in the UNIQUE GAMES instance and where input coordinates of f'_B are permuted according to the constraint between A and B . More precisely,

$$g_A(x, z) = \mathbb{E}_{B \sim G^R(A)} \mathbb{E}_{\tilde{B} \sim T_V(B)} \mathbb{E}_{\pi_B} f'_{\pi_B(\tilde{B})}(\pi_B(x, z)) \quad \text{where } f'_B(x, z) = \mathbb{E}_{(B', x') \sim M_z(B, x)} f(B', x', z).$$

Lemma 6.8. *For every $\tau, \delta > 0$, there exists a constant $c > 0$ such that*

$$\mathbb{E}_F \mathbb{E}_{A \sim V^R} \mathbb{E}_{B \sim G^R(A)} \mathbb{E}_{\tilde{A} \sim T_V(A)} \mathbb{E}_{\tilde{B} \sim T_V(B)} \mathbb{P}_{\pi_A, \pi_B \sim P_{i_k}} \left\{ \pi_A^{-1}(F(\pi_A(\tilde{A}))) = \pi_B^{-1}(F(\pi_B(\tilde{B}))) \right\} > c \mathbb{P}_{A \sim V^R} \{ \exists i. \text{Inf}_i(T_{1-\delta} g_A) > \tau \}.$$

Lemma 6.9. *For every $\nu, \beta, \gamma > 0$, $q \in \mathbb{N}$, and $\rho \in (0, 1)$, there exist $\tau, \delta > 0$ such that*

$$\langle f, Hf \rangle < \Gamma_{\rho^2}(\mathbb{E} f) + 2\gamma + \nu + \beta/\gamma^2 + \mathbb{P}_{A \sim V^R} \{ \exists i. \text{Inf}_i(T_{1-\delta} g_A) > \tau \}.$$

Proof. Recall that Lemma 6.5 shows $\langle f, Hf \rangle = \mathbb{E}_A \|T_{\Omega} g_A\|^2$. The operator T_{Ω} is an R -fold tensor operator with second largest eigenvalue ρ . The invariance principle (Theorem 5.9) asserts that there exist $\tau, \delta > 0$ such that $\|T_{\Omega} g_A\|^2 \leq \Gamma_{\rho^2}(\mathbb{E} g_A) + \nu$ if $\text{Inf}_i(T_{1-\delta} g_A) \leq \tau$ for all coordinates $i \in [R]$. Together with Lemma 6.7, we get

$$\begin{aligned} \langle f, Hf \rangle &= \mathbb{E}_A \|T_{\Omega} g_A\|^2 \leq \Gamma_{\rho^2}(\mathbb{E} f + \gamma \sqrt{\mathbb{E} f}) + \nu + \beta/\gamma^2 + \mathbb{P}_{A \sim V^R} \{ \exists i. \text{Inf}_i(T_{1-\delta} g_A) > \tau \} \\ &\leq \Gamma_{\rho^2}(\mathbb{E} f) + 2\gamma + \nu + \beta/\gamma^2 + \mathbb{P}_{A \sim V^R} \{ \exists i. \text{Inf}_i(T_{1-\delta} g_A) > \tau \} \end{aligned}$$

The second inequality above used that $\Gamma_{\rho^2}(\cdot)$ is 2-Lipschitz. \square

Putting together Lemma 6.8 and Lemma 6.9, we get the following lemma as immediate corollary.

Lemma 6.10. *For every $\beta > 0$, $q \in \mathbb{N}$, and $\rho \in (0, 1)$, there exists $\zeta > 0$ such that either*

$$\langle f, Hf \rangle < \Gamma_{\rho^2}(\mathbb{E} f) + 5\beta^{1/3},$$

or there exists an assignment $F: V^R \rightarrow [R]$ such that the probability (6.2) is at least ζ .

Proof. Choosing $\gamma = \nu = \beta^{1/3}$ in Lemma 6.9, we get that for some $\tau, \delta > 0$,

$$\langle f, Hf \rangle = \mathbb{E}_A \|T_{\Omega} g_A\|^2 \leq \Gamma_{\rho^2}(\mathbb{E} f) + 4\beta^{1/3} + \mathbb{P}_{A \sim V^R} \{ \exists i. \text{Inf}_i(T_{1-\delta} g_A) > \tau \}.$$

Taking $\zeta = c\beta^{1/3}$ for the constant c in Lemma 6.8 then proves the claim. \square

6.2.2 Decoding a very small non-expanding set in G

The following lemma is a slight adaptation of a result in [RS10] (a reduction from SMALL-SET EXPANSION to UNIQUE GAMES). We present a sketch of the proof in Section A.1.

Lemma 6.11. Let G be graph with vertex set V . Let a distribution on pairs of tuples (\tilde{A}, \tilde{B}) be defined by choosing $A \sim V^R$, $B \sim G^{\otimes R}(A)$ and then $\tilde{A} \sim T_V(A)$, $\tilde{B} \sim T_V(B)$. Let $F: V^R \rightarrow [R]$ be a function such that over the choice of random tuples and two random permutations $\pi_A, \pi_B \in \Pi_k$,

$$\mathbb{P}_{(\tilde{A}, \tilde{B})} \mathbb{P}_{\pi_A, \pi_B \sim \Pi_k} \left\{ \pi_A^{-1}(F(\pi_A(\tilde{A}))) = \pi_B^{-1}(F(\pi_B(\tilde{B}))) \right\} \geq \zeta.$$

Then there exists a set $S \subseteq V$ with $\mu(S) \in [\frac{\zeta}{16R}, \frac{3k}{\varepsilon_V R}]$ satisfying $\Phi_G(S) \leq 1 - \frac{\zeta}{16k}$.

Putting together [Lemma 6.11](#) and [Lemma 6.10](#), we get the following lemma — the main lemma for the soundness of the reduction.

Lemma 6.12. Let G be a graph with vertex set V . Let H be the reduction in [Figure 1](#) applied to G with parameters $R, q \in \mathbb{N}$, $\varepsilon_V, \beta > 0$ and $\rho \in (0, 1)$. The vertex set of H is $V^R \times \Omega^R$, where $\Omega = [q] \times \{\perp, \top\}_\beta$. Then there exists $\zeta = \zeta(\beta, q, \rho) > 0$ such that either

$$\forall f: V^R \times \Omega^R \rightarrow [0, 1]. \forall \gamma > 0. \quad \langle f, Hf \rangle < \Gamma_{\rho^2}(\mathbb{E} f) + 5\beta^{1/3},$$

or there exists a vertex set $S \subseteq V$ with $\mu(S) \in [\frac{\zeta}{R}, \frac{3k}{\varepsilon_V R}]$ and $\Phi_G(S) \leq 1 - \zeta/k$.

6.3 Putting things together

Theorem (Restatement of [Theorem 3.5](#)). For all $q \in \mathbb{N}$ and $\varepsilon, \gamma > 0$, it is SSE-hard to distinguish between the following two cases for a given graph $H = (V_H, E_H)$

Yes: There exist q disjoint sets $S_1, \dots, S_q \subseteq V_H$ satisfying for all $l \in [q]$,

$$\mu(S_l) = \frac{1}{q} \quad \text{and} \quad \Phi_H(S_l) \leq \varepsilon + o(\varepsilon).$$

No: For all sets $S \subseteq V_H$,

$$\Phi_H(S) \geq \Phi_{\mathcal{G}(1-\varepsilon/2)}(\mu(S)) - \gamma/\mu(S)$$

where $\Phi_{\mathcal{G}(1-\varepsilon/2)}(\mu(S))$ is the expansion of sets of volume $\mu(S)$ in the infinite Gaussian graph $\mathcal{G}(1-\varepsilon/2)$.

Proof. The follows by proper choice of parameters for the reduction in [Figure 1](#). Given q, ε, γ , we choose the various parameters in the reduction as below:

- $\rho = \sqrt{1 - \frac{\varepsilon}{2}}$.
- $\beta = \min \left\{ \frac{\gamma^3}{200}, \varepsilon \right\}$, so that the error $5\beta^{1/3} < \gamma$ in [Lemma 6.12](#) and the error $\frac{1-\rho^2}{\rho^2}\beta = O(\varepsilon^2)$ in [Lemma 6.2](#).
- $k = \Omega(\log(1/\varepsilon))$, so that the $2^{-\Omega(k)}$ error term in [Lemma 6.2](#) is $O(\varepsilon^2)$.
- $\varepsilon_V = \varepsilon^2$ and $\eta = \min \left\{ \varepsilon^2, \frac{\zeta}{k} \right\}$. Here, $\zeta = \zeta(\beta, q, \rho)$ is the constant given by [Lemma 6.12](#). The above choices ensure that the error term $\varepsilon_V + \eta$ in [Lemma 6.2](#) are $O(\varepsilon^2)$ and $\eta \leq \frac{\zeta}{k}$ for applying [Lemma 6.12](#).
- $M = \max \left\{ \frac{k}{\beta\zeta}, \frac{3\beta}{\varepsilon_V} \right\}$.
- $R = \frac{k}{\beta\delta}$, where $\delta \in (0, \eta)$ is the one for which we intend to show a reduction from SMALL-SET EXPANSION (η, δ, M) .

Given an instance $G = (V, E)$ of SMALL-SET EXPANSION (η, δ, M) , let H be the graph obtained from the reduction in Figure 1 with these parameters.

From Lemma 6.2, we get that the YES case of SMALL-SET EXPANSION (η, δ, M) implies the YES case of the above problem. On the other hand Lemma 6.12 gives that a contradiction to the No case of the above problem produces a set S in G with measure between $\frac{\zeta}{R}$ and $\frac{3k}{\varepsilon VR}$ with $\Phi_G(S) \leq 1 - \frac{\zeta}{k}$. By our choice of parameters, this is a set of measure between $\frac{\delta}{M}$ and $M\delta$ with expansion $1 - \frac{\zeta}{k} \leq 1 - \eta$. This contradicts the No case of SMALL-SET EXPANSION (η, δ, M) . \square

Theorem (Restatement of Theorem 3.2). *For every $q \in \mathbb{N}$ and every $\varepsilon, \gamma > 0$, it is SSE-hard to distinguish between the following cases for a given UNIQUE GAMES instance \mathcal{U} with alphabet size q :*

YES: *The UNIQUE GAMES instance \mathcal{U} is almost satisfiable, $\text{opt}(\mathcal{U}) > 1 - 2\varepsilon - o(\varepsilon)$*

No: *The optimum of the UNIQUE GAMES instance \mathcal{U} is negligible, and the expansion profile of the instance resembles the Gaussian graph $\mathcal{G}(1 - \varepsilon)$. More precisely, the UNIQUE GAMES instance \mathcal{U} satisfies $\text{opt}(\mathcal{U}) < O(q^{-\varepsilon/(2-\varepsilon)}) + \gamma$ and in addition, the constraint graph G of \mathcal{U} satisfies*

$$\forall S \subseteq V(G). \quad \Phi_G(S) \geq \Phi_{\mathcal{G}(1-\varepsilon)}(\mu(S)) - \gamma/\mu(S).$$

Proof. Let all the parameters for the reduction in Figure 1 be chosen as in the proof for Theorem 3.5, replacing ε by 2ε . Let H be the graph generated by the reduction starting from an instance G of SMALL-SET EXPANSION (η, δ, M) . Let \mathcal{U} be the UNIQUE GAMES instance defined on the graph $H/[q]$, as described in Remark 6.1.

We claim that any partition S_1, \dots, S_q of the vertices in H , satisfying the first property in Lemma 6.2, corresponds to an assignment to the vertices in $H/[q]$ and vice-versa. A partition is simply a function $F: V_H \rightarrow [q]$. Restricting the function to the representatives of each equivalence class gives an assignment for the vertices in $H/[q]$. Note that here F also satisfies that $(A, x, z) = [(A, x, z)] + l \implies F((A, x, z)) = F([(A, x, z)]) + l$. Similarly, given an assignment F , we can extend it to all the vertices in H by defining $F((A, x, z)) = F([(A, x, z)]) + l$ if $(A, x, z) = [(A, x, z)] + l$.

In the YES case, we construct an assignment to the UNIQUE GAMES instance \mathcal{U} using the partition S_1, \dots, S_q . The fraction of edges $(\pi_B(B', x'_B, z_B), \pi_C(C', x'_C, z_C))$ that are not satisfied is exactly the probability that $F(\pi_B(B', x'_B, z_B)) \neq F(\pi_C(C', x'_C, z_C))$ for a random edge. However, this is exactly $\mathbb{E}_{l \in [q]} \Phi_H(S_l)$ which is at most $2\varepsilon + o(\varepsilon)$ by Theorem 3.5.

In the No case, we note that we can construct a partition S_1, \dots, S_q from any assignment F . The fraction of unsatisfied edges is again $\mathbb{E}_{l \in [q]} \Phi_H(S_l) \geq 1 - q(\Gamma_{1-\varepsilon}(1/q) + \gamma)$ by Theorem 3.5. Also, any set S in $H/[q]$ corresponds to a set \tilde{S} in H with $\mu(\tilde{S}) = \mu(S)$, where \tilde{S} contains all the vertices for each class in S . The edges leaving S and \tilde{S} are the same and hence their expansion is identical. \square

Acknowledgments

We are grateful to Subhash Khot for suggesting that our techniques should also show that UNIQUE GAMES is SSE-hard on graph with high (small-set) expansion (Theorem 3.2). We also thank Oded Regev for insightful discussions and Boaz Barak for helpful comments on this manuscript.

References

- [ABS10] Sanjeev Arora, Boaz Barak, and David Steurer, *Subexponential algorithms for unique games and related problems*, FOCS, 2010, pp. 563–572.
- [AKK⁺08] Sanjeev Arora, Subhash Khot, Alexandra Kolla, David Steurer, Madhur Tulsiani, and Nisheeth K. Vishnoi, *Unique games on expanding constraint graphs are easy*, STOC, 2008, pp. 21–28.
- [Alo86] Noga Alon, *Eigenvalues and expanders*, Combinatorica **6** (1986), no. 2, 83–96.
- [AM85] Noga Alon and V. D. Milman, λ_1 , *isoperimetric inequalities for graphs, and superconcentrators*, J. Comb. Theory, Ser. B **38** (1985), no. 1, 73–88.
- [AR98] Yonatan Aumann and Yuval Rabani, *An $O(\log k)$ approximate min-cut max-flow theorem and approximation algorithm*, SIAM J. Comput. **27** (1998), no. 1, 291–301.
- [ARV04] Sanjeev Arora, Satish Rao, and Umesh V. Vazirani, *Expander flows, geometric embeddings and graph partitioning*, STOC, 2004, pp. 222–231.
- [BKN⁺10] Nikhil Bansal, Robert Krauthgamer, Viswanath Nagarajan, Konstantin Makarychev, Joseph (Seffi) Naor, and Roy Schwartz, *Min-max graph partitioning and small set expansion*, Manuscript, 2010.
- [Bor85] Christer Borell, *Geometric bounds on the Ornstein-Uhlenbeck velocity process*, Z. Wahrsch. Verw. Gebiete **70** (1985), no. 1, 1–13. MR 795785 (87k:60103)
- [Che70] Jeff Cheeger, *A lower bound for the smallest eigenvalue of the Laplacian*, Problems in analysis (Papers dedicated to Salomon Bochner, 1969), Princeton Univ. Press, Princeton, N. J., 1970, pp. 195–199. MR MR0402831 (53 #6645)
- [DKSV06] Nikhil R. Devanur, Subhash Khot, Rishi Saket, and Nisheeth K. Vishnoi, *Integrality gaps for sparsest cut and minimum linear arrangement problems*, STOC, 2006, pp. 537–546.
- [FS02] Uriel Feige and Christian Scheideler, *Improved bounds for acyclic job shop scheduling*, Combinatorica **22** (2002), no. 3, 361–399.
- [GMPT07] Konstantinos Georgiou, Avner Magen, Toniann Pitassi, and Iannis Tourlakis, *Integrality gaps of $2 - o(1)$ for vertex cover sdps in the Lovász–Schrijver hierarchy*, FOCS, 2007, pp. 702–712.
- [Kho02] Subhash Khot, *On the power of unique 2-prover 1-round games*, STOC, 2002, pp. 767–775.
- [KKMO07] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O’Donnell, *Optimal inapproximability results for MAX-CUT and other 2-variable CSPs?*, SIAM J. Comput. **37** (2007), no. 1, 319–357.
- [KR08] Subhash Khot and Oded Regev, *Vertex cover might be hard to approximate to within $2 - \epsilon$* , J. Comput. Syst. Sci. **74** (2008), no. 3, 335–349.
- [KS09] Subhash Khot and Rishi Saket, *SDP integrality gaps with local ℓ_1 -embeddability*, FOCS, 2009.

- [KV05] Subhash Khot and Nisheeth K. Vishnoi, *The unique games conjecture, integrality gap for cut problems and embeddability of negative type metrics into ℓ_1* , FOCS, 2005, pp. 53–62.
- [LLR95] Nathan Linial, Eran London, and Yuri Rabinovich, *The geometry of graphs and some of its algorithmic applications*, Combinatorica **15** (1995), no. 2, 215–245.
- [LR99] Frank Thomson Leighton and Satish Rao, *Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms*, J. ACM **46** (1999), no. 6, 787–832.
- [MOO05] Elchanan Mossel, Ryan O’Donnell, and Krzysztof Oleszkiewicz, *Noise stability of functions with low influences invariance and optimality*, FOCS, 2005, pp. 21–30.
- [Räc08] Harald Räcke, *Optimal hierarchical decompositions for congestion minimization in networks*, STOC, 2008, pp. 255–264.
- [Rag08] Prasad Raghavendra, *Optimal algorithms and inapproximability results for every CSP?*, STOC, 2008, pp. 245–254.
- [RS09] Prasad Raghavendra and David Steurer, *Integrality gaps for strong SDP relaxations of unique games*, FOCS, 2009, pp. 575–585.
- [RS10] ———, *Graph expansion and the unique games conjecture*, STOC, 2010, pp. 755–764.
- [RST10] Prasad Raghavendra, David Steurer, and Prasad Tetali, *Approximations for the isoperimetric and spectral profile of graphs and related parameters*, STOC, 2010, pp. 631–640.

A Further Proofs

A.1 Reduction from SMALL-SET EXPANSION to UNIQUE GAMES

In this section, we sketch a proof of the following slight adaption of a result in [RS10].

Lemma (Restatement of Lemma 6.11). *Let G be graph with vertex set V . Let a distribution on pairs of tuples (\tilde{A}, \tilde{B}) be defined by choosing $A \sim V^R$, $B \sim G^{\otimes R}(A)$ and then $\tilde{A} \sim T_V(A)$, $\tilde{B} \sim T_V(B)$. Let $F: V^R \rightarrow [R]$ be a function such that over the choice of random tuples and two random permutations $\pi_A, \pi_B \in \Pi_k$,*

$$\mathbb{P}_{(\tilde{A}, \tilde{B})} \mathbb{P}_{\pi_A, \pi_B \sim \Pi_k} \left\{ \pi_A^{-1}(F(\pi_A(\tilde{A}))) = \pi_B^{-1}(F(\pi_B(\tilde{B}))) \right\} \geq \zeta.$$

Then there exists a set $S \subseteq V$ with $\mu(S) \in \left[\frac{\zeta}{16k}, \frac{3k}{\varepsilon V R} \right]$ satisfying $\Phi_G(S) \leq 1 - \frac{\zeta}{16k}$.

Let $R' = R/k$ and let $\tilde{A}_{R'}, \tilde{B}_{R'}$ denote tuples of length R' generated by a process similar to the used for generating \tilde{A}, \tilde{B} (which have length R). Using the reduction from partial to total unique games in [RS10], we can show the following for completely random permutations (instead of block-wise random) permutations $\pi'_A, \pi'_B: [R'] \rightarrow [R']$.

Claim A.1. *Given a function $F: V^R \rightarrow [R]$ as above, there exists a function $F': V^{R'} \rightarrow [R']$ such that*

$$\mathbb{P}_{(\tilde{A}_{R'}, \tilde{B}_{R'})} \mathbb{P}_{\pi'_A, \pi'_B \sim S_{R'}} \left\{ \pi'_A^{-1}(F'(\pi'_A(\tilde{A}_{R'}))) = \pi'_B^{-1}(F'(\pi'_B(\tilde{B}_{R'}))) \right\} \geq \zeta/k.$$

Proof. We construct a randomized function F' which given an R' -tuple, embeds it as one of the blocks (of size R') in a random R -tuple, and then outputs a value according to the value of F on the R -tuple.

Formally, let $\tilde{A}_{R-R'}, \tilde{B}_{R-R'}$ denote tuples of size $R - R'$ generated by independently picking each pair of coordinates to be an edge in G with noise ε_V . For $j \in [k]$, let $\tilde{A}_{R-R'} +_j \tilde{A}_{R'}$ denote an R -tuple generated by inserting $\tilde{A}_{R'}$ after the $(j-1)^{th}$ block in $\tilde{A}_{R-R'}$. Let $\pi_A^{(j)}$ be a random permutation in Π_k which is equal to identity on the j^{th} block. To define F' , we then generate a random j , $\tilde{A}_{R-R'}, \tilde{B}_{R-R'}$ and fix it *globally*. For each input $\tilde{A}_{R'}$, we then *independently* choose $T_A \in \{\tilde{A}_{R-R'}, \tilde{B}_{R-R'}\}$, a permutation $\pi_A^{(j)}$, and define $F'(\tilde{A}_{R'})$ as:

$$F'(\tilde{A}_{R'}) = \begin{cases} (\pi_A^{(j)})^{-1} \left(F \left(\pi_A^{(j)}(T_A) +_j \tilde{A}_{R'} \right) \right) - (j-1)R/k & (\pi_A^{(j)})^{-1} \left(F \left(\pi_A^{(j)}(T_A) +_j \tilde{A}_{R'} \right) \right) \in ((j-1)R/k, jR/k) \\ 1 & \text{otherwise} \end{cases}$$

Let π_A be the permutation which is $\pi_A^{(j)}$ on the j^{th} block and $\pi_A^{(j)}$ elsewhere. Define π_B similarly. Note that both π_A, π_B are distributed as random elements of Π_k . Conditioned on $T_A = \tilde{A}_{R-R'}, T_B = \tilde{B}_{R-R'}$ (or vice-versa), the required probability is at least

$$\begin{aligned} & \mathbb{P}_{j \in [k]} \mathbb{P}_{(\tilde{A}_{R-R'}, \tilde{B}_{R-R'})} \mathbb{P}_{(\tilde{A}_{R'}, \tilde{B}_{R'})} \mathbb{P}_{\pi_A, \pi_B \sim \Pi_k} \left\{ \pi_A^{-1}(F(\pi_A(\tilde{A}_{R-R'} +_j \tilde{A}_{R'}))) = \pi_B^{-1}(F(\pi_B(\tilde{B}_{R-R'} +_j \tilde{B}_{R'}))) \right\} \\ &= \frac{1}{k} \cdot \mathbb{P}_{(\tilde{A}, \tilde{B})} \mathbb{P}_{\pi_A, \pi_B \sim \Pi_k} \left\{ \pi_A^{-1}(F(\pi_A(\tilde{A}))) = \pi_B^{-1}(F(\pi_B(\tilde{B}))) \right\} = \frac{\zeta}{k}. \end{aligned}$$

Since we have $T_A = \tilde{A}_{R-R'}, T_B = \tilde{B}_{R-R'}$ or vice-versa with probability $1/2$, the required probability is at least $\zeta/2k$. \square

Let ζ' denote $\zeta/2k$.

To construct the set $S \subseteq V$, we proceed as in [RS10] by defining the influence of a single vertex on the output of F . For $U \in V^{R'-1}$ and $v \in V$, let $\tilde{U} \sim T_V(U)$ and $\tilde{v} \sim T_V(v)$. For $i \in R'$, we use $\tilde{U} +_i \tilde{v}$ to denote the tuple $(\tilde{U}_1, \dots, \tilde{U}_{i-1}, v, \tilde{U}_i, \dots, \tilde{U}_{R'}) \in V^{R'}$ in which v is inserted at the i th position. We define the function $F_U(v)$, which measures how often is the index of v chosen by F , when applied to a random permutation π of $\tilde{U} +_i \tilde{v}$.

$$F_U(v) := \mathbb{E}_{\tilde{U} \sim T_V(U)} \mathbb{E}_{\tilde{v} \sim T_V(v)} \mathbb{P}_{i \in [R']} \mathbb{P}_{\pi \in S_{R'}} \left\{ F(\pi(\tilde{U} +_i \tilde{v})) = \pi(i) \right\}$$

We shall need the following (slight variants of) statements proved in [RS10]. We include the proofs in the appendix for completeness.

Lemma A.2 (Glorified Markov Inequality). *Let Ω be a probability space and let $X, Y: \Omega \rightarrow \mathbb{R}_+$ be two jointly distributed non-negative random variables over Ω . Suppose $\mathbb{E} X \leq \gamma \mathbb{E} Y$. Then, there exists $\omega \in \Omega$ such that $X(\omega) \leq 2\gamma Y(\omega)$ and $Y(\omega) \geq \mathbb{E} Y/2$.*

Proposition A.3. *Let $F: V^{R'} \rightarrow [R']$ satisfy $\mathbb{P}_{(\tilde{A}, \tilde{B})} \mathbb{P}_{\pi_A, \pi_B} \left\{ \pi_A^{-1}(F(\pi_A(\tilde{A}))) = \pi_B^{-1}(F(\pi_B(\tilde{B}))) \right\} \geq \zeta'$, and the functions $F_U: V \rightarrow [0, 1]$ be defined as above. Then,*

1.

$$\mathbb{E}_{U \sim V^{R'-1}} \mathbb{E}_{v \sim V} F_U(v) = \frac{1}{R'}, \quad (\text{A.1})$$

2. *for all $U \in V^{R'-1}$,*

$$\mathbb{E}_{v \sim V} F_U(v) \leq \frac{2}{\varepsilon_V R'}, \quad (\text{A.2})$$

3.

$$\mathbb{E}_{(U,W) \in E^{R'-1}} \mathbb{E}_{(v_1,v_2) \in E} [F_U(v_1)F_W(v_2)] \geq \frac{\zeta'}{R'}. \quad (\text{A.3})$$

Assuming Lemma A.2 and Proposition A.3, we can now complete the proof of Lemma 6.11.

Proof. By (A.3) and (A.1), we have that

$$\mathbb{E}_{(U,W) \in E^{R'-1}} \mathbb{E}_{(v_1,v_2) \in E} [F_U(v_1)F_W(v_2)] \geq \frac{\zeta'}{R'} = \frac{\zeta'}{2} \cdot \mathbb{E}_{(U,W) \in E^{R'-1}} \mathbb{E}_{(v_1,v_2) \in E} [F_U(v_1) + F_W(v_2)]$$

Using Lemma A.2, this gives that there exist $(U^*, W^*) \in E^{R'-1}$ such that

$$\mathbb{E}_{(v_1,v_2) \in E} [F_{U^*}(v_1)F_{W^*}(v_2)] \geq \frac{\zeta'}{2R'} \quad \text{and} \quad \mathbb{E}_{(v_1,v_2) \in E} [F_{U^*}(v_1)F_{W^*}(v_2)] \geq \frac{\zeta'}{4} \cdot (\mathbb{E} F_{U^*} + \mathbb{E} F_{W^*})$$

We now construct the set S randomly, by choosing each $v \in V$ to be in S with probability $(F_{U^*}(v) + F_{W^*}(v))/2$. We first check that the expected volume of the set is large.

$$\begin{aligned} \mathbb{E} \mu(S) &= \mathbb{E}_{v \sim V} \left[\frac{F_{U^*}(v) + F_{W^*}(v)}{2} \right] = \mathbb{E}_{(v_1,v_2) \in E} \left[\frac{F_{U^*}(v_1) + F_{W^*}(v_2)}{2} \right] \\ &\geq \mathbb{E}_{(v_1,v_2) \in E} \left[\frac{F_{U^*}(v_1)F_{W^*}(v_2)}{2} \right] \quad (\text{Using } a + b \geq ab \text{ for } a, b \in [0, 1]) \\ &\geq \frac{\zeta'}{4R'} \end{aligned}$$

Combining this with (A.2), we get that $\mathbb{E} \mu(S) \in \left[\frac{\zeta'}{4R'}, \frac{2}{\varepsilon_V R'} \right]$. Also, by a Chernoff bound, we have that with probability $1 - \exp(-\Omega(|V|))$, $\mu(S) \in \left[\frac{\zeta'}{8R'}, \frac{3}{\varepsilon_V R'} \right]$.

To show that the expansion of the set is bounded away from 1, we show a lower bound on the expected number of edges that stay within the set, denoted by $G(S, S)$.

$$\begin{aligned} \mathbb{E} G(S, S) &= \mathbb{E}_{(v_1,v_2) \in E} \left[\left(\frac{F_{U^*}(v_1) + F_{W^*}(v_1)}{2} \right) \left(\frac{F_{U^*}(v_2) + F_{W^*}(v_2)}{2} \right) \right] \\ &= \frac{1}{2} \cdot \mathbb{E}_{(v_1,v_2) \in E} F_{U^*}(v_1)F_{W^*}(v_2) + \frac{1}{4} \cdot \mathbb{E}_{(v_1,v_2) \in E} [F_{U^*}^2(v_1) + F_{W^*}^2(v_2)] \\ &\geq \frac{1}{2} \cdot \mathbb{E}_{(v_1,v_2) \in E} F_{U^*}(v_1)F_{W^*}(v_2) \\ &\geq \frac{\zeta'}{4} \mathbb{E}_{v \sim V} [(F_{U^*} + F_{W^*})/2] = \frac{\zeta'}{4} \cdot \mathbb{E} \mu(S) \end{aligned}$$

Thus, we have

$$\mathbb{E} [G(S, S) - \frac{\zeta'}{8} \mu(S)] \geq \frac{\zeta'}{8} \mathbb{E} \mu(S) \geq \frac{\zeta'^2}{8R'}.$$

In particular, we get that with probability at least $\frac{\zeta'^2}{16R'}$ over the choice of S , $G(S, S) \geq \frac{\zeta'}{8} \cdot \mu(S)$. Hence, with probability $\frac{\zeta'^2}{16R'} - e^{-\Omega(|V|)}$, we have $\mu(S) \in \left[\frac{\zeta'}{8R'}, \frac{3}{\varepsilon_V R'} \right]$ and $G(S, S) \geq \frac{\zeta'}{8} \cdot \mu(S)$. For such a set we have $\Phi_G(S) = 1 - (G(S, S)/\mu(S)) \leq 1 - \frac{\zeta'}{8}$, which proves the claim. \square

A.2 Stronger Small-Set Expansion Hypothesis

Proposition (Restatement of [Proposition 5.8](#)). *For all $\eta > 0$, $M \geq 1$ and all $\delta < 1/M$, there is polynomial time reduction from SMALL-SET EXPANSION $(\frac{\eta}{M}, \delta)$ to SMALL-SET EXPANSION (η, δ, M) .*

Proof. Let $\eta' = \frac{\eta}{M}$. The reduction is in fact, the trivial one which, given an instance $G = (V, E)$ of SMALL-SET EXPANSION (η', δ) treats as an instance of SMALL-SET EXPANSION (η, δ, M) . If we are in the YES case of SMALL-SET EXPANSION (η', δ) , then there is a set S with $\mu(S) = \delta$ and $\Phi_G(S) \leq \eta' \leq \eta$. Hence, we are also in the YES case of SMALL-SET EXPANSION (η, δ, M) .

For the other direction, assume that we are *not* in the No case of SMALL-SET EXPANSION (η, δ, M) and there exists a set S with $\mu(S) \in (\frac{\delta}{M}, \frac{\delta}{M})$ and $\Phi_G(S) \leq 1 - \eta$. Then the fraction of edges $G(S, S)$ stay inside S is at least $\eta \cdot \mu(S)$. If $\mu(S) \geq \delta$, then we randomly sample a subset S' of S with volume δ . For each edge $(u, v) \subseteq S$, the chance that $(u, v) \in S'$ is $\delta^2/\mu(S)^2$. Then

$$\mathbb{E} \Phi_G(S') = 1 - \frac{\mathbb{E} G(S', S')}{\delta} \leq 1 - \frac{(\delta^2/\mu(S)^2) \cdot \eta \cdot \mu(S)}{\delta} \leq 1 - \frac{\eta}{M}.$$

Then, we cannot be in the No case of SMALL-SET EXPANSION (η', δ) . When $\mu(S) \leq \delta$, we simply create a set S' by adding extra vertices to S to increase its measure to δ . Then,

$$\Phi_G(S') = 1 - \frac{G(S', S')}{\delta} \leq 1 - \frac{G(S, S)}{\delta} \leq 1 - \frac{\eta \cdot \mu(S)}{\delta} \leq 1 - \frac{\eta}{M}.$$

□

A.3 Hardness of MINIMUM LINEAR ARRANGEMENT and BALANCED SEPARATOR

Corollary A.4 (Hardness of BALANCED SEPARATOR and MIN BISECTION). *There is a constant c such that for arbitrarily small $\varepsilon > 0$, it is SSE-hard to distinguish the following two cases for a given graph $G = (V, E)$:*

YES: *There exists a cut $(S, V \setminus S)$ in G such that $\mu(S) = \frac{1}{2}$ and $\Phi_G(S) \leq \varepsilon + o(\varepsilon)$.*

NO: *Every cut $(S, V \setminus S)$ in G , with $\mu(S) \in (\frac{1}{10}, \frac{1}{2})$ satisfies $\mu_G(S) \geq c\sqrt{\varepsilon}$.*

Proof. The result follows immediately by applying [Theorem 3.5](#) with the given ε and taking $q = 2$, $\gamma = o(\sqrt{\varepsilon})$. In the No case we get that for all sets S with $\mu(S) \in (\frac{1}{10}, \frac{1}{2})$, $G(S, S) \leq \Gamma_{1-\varepsilon/2}(1/10) + o(\sqrt{\varepsilon}) \leq \mu(S)(1 - c\sqrt{\varepsilon}) + o(\sqrt{\varepsilon})$ for some $c > 0$. Thus $\Phi_G(S) \geq c'\sqrt{\varepsilon}$ for some $c' > 0$. □

The following corollary uses the fact that in the YES case of [Theorem 3.5](#), we actually partition the graph into many non-expanding sets instead of finding just one such set.

Corollary A.5 (Hardness of MINIMUM LINEAR ARRANGEMENT). *It is SSE-hard to approximate MINIMUM LINEAR ARRANGEMENT to any fixed constant factor. Formally, there exists $c > 0$ such that for every $\varepsilon > 0$, it is SSE-hard to distinguish between the following two cases for a given graph $G = (V, E)$, with $|V| = n$:*

YES: *There exists an ordering $\pi: V \rightarrow [n]$ of the vertices such that $\mathbb{E}_{(u,v) \sim E} [|\pi(u) - \pi(v)|] \leq \varepsilon n$*

NO: *For all orderings $\pi: V \rightarrow [n]$, $\mathbb{E}_{(u,v) \sim E} [|\pi(u) - \pi(v)|] \geq c\sqrt{\varepsilon}n$*

Proof. Apply [Theorem 3.5](#) taking $q = \lceil 2/\varepsilon \rceil$, $\varepsilon' = \varepsilon/3$ and $\gamma = \varepsilon$. In the YES case, we pick an arbitrary ordering π which orders elements in each of the sets S_1, \dots, S_q contiguously. For these sets, all edges in the set have length at most n/q and at most $\varepsilon' + o(\varepsilon)$ fraction of the edges leave the sets. Thus,

$$\mathbb{E}_{(u,v) \sim E} [|\pi(u) - \pi(v)|] \leq \frac{n}{q} + \varepsilon' n + o(\varepsilon n) \leq \varepsilon n$$

The proof for the No case follows from an observation of [\[DKSV06\]](#), that for a graph G if every set S with $\mu(S) \in (\frac{1}{3}, \frac{1}{2})$ has $G(S, V \setminus S) \geq \theta$, then for any ordering $\pi: V \rightarrow [n]$, $\mathbb{E}_{(u,v) \sim E} [|\pi(u) - \pi(v)|] \geq \frac{\theta}{3} \cdot n$ (else one can obtain a contradiction by optimally ordering the points and cutting randomly between the positions $n/3$ and $2n/3$). Here, $G(S, V \setminus S) \geq 1/3 - \Gamma_{1-\varepsilon/6}(1/3) \geq c' \sqrt{\varepsilon}$ for some $c' > 0$. Thus, $\mathbb{E}_{(u,v) \sim E} [|\pi(u) - \pi(v)|] \geq \frac{c'}{3} \cdot n$ for all $\pi: V \rightarrow [n]$. \square